

For Reference

NOT TO BE TAKEN FROM THIS ROOM

Ex libris
UNIVERSITATIS
ALBERTAENSIS



THE UNIVERSITY OF ALBERTA

RELEASE FORM

NAME OF AUTHOR . . . Jerome Edmond Bickenbach

TITLE OF THESIS . . . The Nature and Scope
of Reflective Reasoning

DEGREE FOR WHICH THESIS WAS PRESENTED . . . Doctor of Philosophy . . .

YEAR THIS DEGREE GRANTED . . . 1977

Permission is hereby granted to THE UNIVERSITY OF ALBERTA LIBRARY
to reproduce single copies of this thesis and to lend or sell such copies
for private, scholarly or scientific research purposes only.

The author reserves other publication rights, and neither the thesis nor extensive extracts from it may be printed or otherwise reproduced without the author's written permission.

THE UNIVERSITY OF ALBERTA

THE NATURE AND SCOPE OF REFLECTIVE REASONING

by



JEROME EDMOND BICKENBACH

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES AND RESEARCH


IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

DEPARTMENT.....PHILOSOPHY.....

EDMONTON, ALBERTA

FALL, 1977



Digitized by the Internet Archive
in 2024 with funding from
University of Alberta Library

<https://archive.org/details/Bickenbach1977>

THE UNIVERSITY OF ALBERTA
FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research, for acceptance, a thesis entitled THE NATURE AND SCOPE OF REFLECTIVE REASONING submitted by Jerome Edmond Bickenbach, in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Philosophy.

ABSTRACT

In this thesis, I attempt to provide the theoretical framework for a conception of reasoning which, though neither deductive nor inductive in form, is both legitimate and suitable for questions requiring reflection in ethics, philosophy, the law and elsewhere. I begin by isolating a particularly prevalent dogma about the nature of reasoning, justification and proof, Deductivism--the view that, in the end, only deduction is legitimate reasoning, only the deductive model is appropriate to justification and proof. I next consider deductive reasoning as reasoning from, and to, principles and suggest that there is a genuine epistemological dilemma which applies to this form of reasoning. I then suggest that non-deductive reasoning from particulars--case-by-case reasoning--can save deduction from this dilemma, but only if the legitimacy of this non-deductive form of reasoning is acknowledged. After showing that the kind of reasoning I have in mind is non-inductive as well as non-deductive, I attempt to provide grounds for the legitimacy of reflective reasoning. I conclude by considering the kinds of question for which reflective reasoning is suitable.

CONTENTS

Chapter

I. INTRODUCTION	1
Deductivism	1
Deductivism in Ethics	17
II. REASONING FROM PRINCIPLES	30
A Dilemma for Deduction--First Formulation	31
A Dilemma for Deduction--Second Formulation	44
Reasoning from Principles	58
Footnotes to Chapter II	76
III. REASONING FROM PARTICULARS	79
The 'End' of Justification	80
Reasoning from Particulars and the Law	89
Questions about Concepts	115
Footnotes to Chapter III	137
IV. EXPECTATION REASONING	139
Induction--Utility and Legitimacy	140
Induction and Reflection	164
Footnotes to Chapter IV	197
V. REFLECTIVE REASONING	199
Bringing the Pieces Together	199
Analogical Reasoning	215
Rhetoric and Argumentation	232
Reflection and Legitimacy	257
Footnotes to Chapter V	300
VI. QUESTIONS REQUIRING REFLECTION	302
.	
BIBLIOGRAPHY	328

CHAPTER I

INTRODUCTION

Deductivism

Philosophical problems and disputes are often the result of philosophers asking peculiar questions. These questions seem to possess a perplexing conjunction of properties. They are questions we are convinced are answerable--indeed questions whose answers we may very well be convinced of--but which do not seem to be answerable either deductively or inductively from premisses or facts available to us. If the source of philosophising is wonder, then typically the cause of this wonder is a peculiar, perplexing, bewildering question: How do we know that minds other than our own exist? How can I be sure that memory does not always deceive me? Is it ever possible to prove statements about what is morally right or wrong? Do we really have adequate grounds for thinking that the future will be like the past? Do people ever do things for other than selfish motives? Could it be that everyone is insane? Do I know for certain that I have never been on the surface of Jupiter? These are all peculiar questions, questions which have taxed the philosophical imagination.

Approaching the question of other minds or memory or proof in morality the philosopher assumes a particular stance: his job is to look deeper into the question, to give reasons for being sure which are better reasons than those that come immediately to mind. No one could seriously doubt that other minds exist--that seems clear enough. But the philosopher requires more than that sort of assurance. Traditionally,

the search for better grounds for thinking that other minds exist, that memory is generally reliable or that some moral assertions are provable has been directed toward more general knowledge claims which, were they secure, would deductively guarantee the truth of other knowledge claims. Memory is generally reliable, it could be argued, because it is impossible that God would deceive me by planting phony memories in my mind. But this provides only temporary relief: the philosopher must go on to support, to make secure, the more general knowledge claim. We would have conclusive reasons for thinking that other minds exist if whenever bodies behave in certain ways there are minds at work. But then, how do we support that claim. Or, again, some moral assertions could be proved if they followed deductively from moral principles or axioms about which no doubts could be raised. But are there such principles, and if there are how can we be sure that they are correct?

There would appear to be something about these philosophical questions which makes deduction an ineffectual rational tool: the problems they raise are only pushed further back by the attempt to resolve them deductively. But reasoning inductively, we never seem able to come to grips with the philosophical perplexity. It doesn't seem to be enough to say that so far the earth has completed its journey around the sun on a regular basis. It doesn't seem to be enough to say that other human bodies behave in ways akin to ways I have behaved in the light of my mental life. The questions which trouble the philosopher are questions which seem to stand in need of resolution, not empirical evidence.

Because these philosophical questions do seem answerable, and indeed seem immediately answerable, some philosophers have suggested that we must have a special way of knowing the truth of the commonsense answers we have for them, special modes of knowing which are both immediate and infallible.

These philosophers have, in effect, tried to come to the aid of deduction. They have argued that when in our search for the justifying knowledge claim, the general claim from which we could deduce the truth of other claims, we arrive at something which is ultimately general--the non-deceptive character of God, the indubitability of sense-data reports or the Principle of Utility--at this point intuition steps in. This philosophical move, however, has not been viewed to be an entirely satisfactory one. Of course, we do know that there are other minds; but the philosophical question is not one which is answered by bringing that piece of information to the forefront, underscored, as it were, by the notion of intuition. The philosophical problem of other minds is perplexing because it is asked in the light of what everyone knows and no one would doubt.

There is yet another tactic that has been adopted, however. It has been claimed that these questions are not really questions at all; they seem to be questions--they have the correct grammatical form of a question--but genuine questions are not like these. In the heyday of Logical Empiricism the notions of a 'pseudo-question' and a 'pseudo-problem' were introduced, and introduced in the spirit of revolution. Reading A. J. Ayer's Language, Truth and Logic one cannot help but feel the sense of triumph which this particular school of thought expressed: not everything that looks like a question need be one, and philosophers need no longer be bothered by questions which cannot be answered by means either of scientific investigation or rigorous, demonstrative proof. The sphere of human knowledge can be neatly divided into the analytical truths of logic and mathematics and the synthetic truths of natural science--the rest, the metaphysical remainder, must be swept away. Questions which cannot be answered by deductive or inductive

means are no longer perplexing, for they are no longer legitimate questions: they fall below the line of demarcation between what is sense and what is nonsense.

But the reform was shortlived, the revolution failed. Philosophical questions of this perplexing sort cannot be dealt with so easily, by fiat. Even if they are peculiar, still they have a hold on us, in reflective moments they still intrigue. And they intrigue primarily because, in the context in which they are considered, it is clear what fails to count as an answer. When the philosopher asks whether we ever truly remember things he is asking a question which should strike us in a different way than the question, Do you remember where you left the car keys? The philosophical question is directed to the issue of grounds and warrants; it is not resolved by evidence that might be offered to show that what someone claimed to have remembered was indeed the case. It is not a question which goes away that easily. It is rather a question which is intended to be asked and considered in the light of all the evidence we could possibly need. The philosophical questions which have given rise to various accounts and theories, in short, continued to hold a fascination.

But, more importantly, these questions could not be easily shelved because they bear an affinity to countless other questions which, though not particularly philosophical, seem also to possess the same perplexing conjunction of properties when asked in particular circumstances. The scientist may ponder the question of whether light is particulated, or whether a virus is really alive--and he may ponder these questions after the facts are in, the theories constructed. The judge may be concerned about whether a particular act is a case of negligence or not: the facts are in, the two counsels have made their cases, but ought we really to

say that Jones was negligent? Here deduction may be ineffectual. It may very well be true that every case of failure to exercise reasonable care is a case of negligence; but this fact will not help the judge in his deliberation if it is not clear whether Jones failed to exercise reasonable care. And how do we determine when an action is rude, silly, contrived, heroic, cowardly; are there strict rules which inform us when an action is malicious, cruel, playful, graceful, unseemly, polite, manly, childish, friendly, offensive, audacious, imitative or spontaneous? On the face of it these latter, everyday, kinds of questions are questions about facts, questions about what is right before our eyes. But there are occasions when what is before our eyes is not exactly silly, nor exactly serious, but something on the borderline between the two. Even when we cannot tell whether a particular action is silly, the question whether it is or not is one which we are convinced is answerable, though not by means of deductive or inductive reasoning. At times we are with mundane questions in the same situation we are in with extramundane philosophical questions.

It is, of course, a bit shocking that everyday questions about silliness and rudeness seem to be on all fours with the deep and troubling philosophical questions which the Logical Empiricists wanted to sweep away. We do seem to know when people are sad and actions are cruel or playful; but do we know these things as a result of deductive or inductive reasoning? It would be at least peculiar to say that that is how we know these things. We learn to recognise sadness, cruelty and silliness, we know what these things are. Experience is required, to be sure. But, surely, we do not deduce Smith's cruelty from a general principle which tells us what cruelty is, in all of its possible

manifestations. And do we appeal to induction when we claim that an up-turned cardboard box which children are playing cards on is, on this occasion, a table? We recognise it as such, we may even have to think the matter over a bit; but we do not inductively reason to that conclusion. The Logical Empiricist would not be willing to dismiss as 'pseudo' and nonsensical all questions for which deductive and inductive reasoning seem inappropriate, so why should philosophical questions with this property be suspect?

Friedrich Waismann once characterised a philosopher as a "man who senses as it were hidden crevices in the build of our concepts where others only see the smooth path of commonplaceness before him" (2). The philosopher is someone who worries about things, properties, about minds, beliefs, rights, wrongs and gods. He is troubled about what it means to say that we see a tree over there, that the sun will rise tomorrow because it always has in the past, that numbers exist, people think and human fetuses are persons. He is troubled by concepts, primarily. Sometimes he is concerned about concepts which are the building blocks of our commonsense understanding of the world; sometimes more exotic concepts trouble him. But perplexing metaphysical questions are questions about concepts, questions which may point to crevices in the build of our commonsense concepts: Do we really know that other people see the same colours we do?

To say that the crevices which the philosopher detects are not really there is to say that his questions and problems are pseudo-questions and pseudo-problems. But, to say that the crevices are not really there is also to suppose that all concepts are like those which form the building blocks of logic and mathematics. The logician's

concepts--e.g. well-formedness, deducibility, theorem, and proof--are explicitly defined; the boundaries of his concepts are precisely set out. There are no borderline cases, no crevices to point out and question. If something is a well-formed expression in a formal deductive calculus, it can be shown rigorously to be so: How do we know that ____ is a well-formed expression in this language? is a question with an answer, and it is a question whose answer can be conclusively demonstrated to anyone who cares to listen. But not every concept can be treated in this fashion; some concepts resist such a treatment. There are, as a consequence, questions about concepts, ordinary--Is a mobile home on a foundation still a vehicle?--and specialised--Is angina pectoris a disease or a symptom of a disease? Philosophical questions, I want to suggest, bear a closer relationship to the kinds of questions that are raised in a court of law than they do to the kinds of questions answered by the machinery of logic or the empirical investigations of a scientist.

The view has developed out of some remarks made by Ludwig Wittgenstein that questions about concepts are within the scope of reason despite the fact that neither deductive nor inductive reasoning seems appropriate to them. This view, for the most part, has been developed by John Wisdom who has written extensively on the notion of 'case-by-case' reasoning. It is Wisdom's contention that there are rationally legitimate ways of reasoning that are neither deductive nor inductive but are nonetheless ways of approaching, if not always resolving, philosophically perplexing questions as well as questions concerning concepts which require reasoning in other areas--most notably, the law. The central project of this thesis will be to set out a case in favour of this view. I shall try to show why the sort of reasoning Wisdom has

in mind counts as reasoning, and why it is in an important sense prior to both deductive and inductive reasoning. Along the way I hope to shed some light on the nature of both deduction and induction as forms of reasoning, their strengths and their weaknesses, and the source of their rationality. My principal goal, however, is to provide at least part of a theoretical account of this sort of reasoning--reasoning which I wish to call reflective reasoning, or reflection--by means of an explanation of the kinds of questions which stand in need of reflection.

Now the view I wish to build a case for comes in direct conflict with a very general and entrenched philosophical dogma. It is a dogma about the nature of reasoning, a philosophical position of far-reaching influence which is in effect an attempt to neutralise the perplexity caused by philosophical questions about concepts. I want to call this dogma Deductivism.

Deductivism is the view, first off, that there are only two forms of reasoning: deduction and induction, the first ubiquitous and stable, ultimately secure and unproblematic, the second, although philosophically troublesome (because non-deductive), essential for certain kinds of questions. Deductivism, moreover, is the view that the deductive model of reasoning is the only form suitable for resolving a priori questions, and is, therefore, the paradigm of philosophical reasoning. Deductivism is the dogma--built as it is on Hume's strict division between reasoning about 'relations of ideas' and about 'matters of fact'--that a piece of reasoning must either fit the deductive mould or, failing that, be justifiable in the light of new information, non-inferentially obtained. Hence, on this view, if a question cannot be answered directly by bringing forward new information--new data, new factual information directly obtained by observation--then it is a question which, if it

is to be answered at all, must be answered deductively. But, if all the information is available, and if deduction is for one reason or another inapplicable then, the dogma holds, whatever question remains is either a mere question of words ('Well, call it a table if you like') or else permanently irresolvable ('There is just no way of knowing whether Jones was negligent or not'). As a dogma about reasoning, Deductivism is, ipso facto, also a dogma about a fairly extensive family of concepts. It is a dogma about what it means to be rational, to argue in a rational manner and to arrive at decisions rationally. It is a dogma about what an argument must be like. It is a dogma about what constitutes a proof, what counts as justification and what counts as an explanation. It is, finally, a dogma about the nature of legitimate, arguable, resolvable questions.

It is my contention that Deductivism is a view which most philosophers bring to philosophical inquiry, a view presupposed by philosophers as they attempt to tackle the central questions of various branches of philosophy. The effects of this imposition of Deductivism are extensive, influencing both substantive and methodological issues. As an illustrative example, I shall be setting out in some detail the consequences of the assumption of Deductivism in the case of ethics in the next section. For the remainder of this section, since these modes of reasoning will be our concern in what follows, I want to isolate some of the salient features of deduction and induction.

To begin with, in future I shall mean by an 'inference' an isolatable deductive or inductive 'move', a move from a set of premisses (possibly empty) to a conclusion which is justified by means of a rule of inference (e.g. modus ponens or 'From: m/n of all observed P's have been Q's; infer: m/n of all P's are Q's'). Since 'inference' is, generally

speaking, a term which closely aligned with deductive and inductive logic, rather than cause terminological confusion by speaking of 'non-deductive and non-inductive inferences', in what follows I want to restrict the term to deductive and inductive moves. The general class of which deductive and inductive inferences form a part will be called the class of 'argumentative moves'. In the same spirit, I will use the term 'logic' to mean a deductive or an inductive logic, where a logic might be roughly characterised as a system of rules of inference. A formal logic or logistic system, on the other hand, consists of a language (typically symbolic) with clearly demarked elements (e.g. propositional constants, logical connectives, terms, variables and so on) and an associated syntax which explicitly defines well-formedness, sentencehood and proof in the system, the latter in terms of a given set of inference rules and, possibly, a set of axioms. Whenever formal systems are the object of discussion, I shall assume a standard semantics for the type of formal logic being considered.

One last key notion should be mentioned, the notion of an argument. I am not at all interested in a precise and overly restrictive characterisation of an argument. This particular notion I want to keep as open-ended as possible at this stage. Stuart Hampshire has offered a suitably versatile characterisation:

If one criticises a sequence of sentences by saying that assertion or denial of the earlier members of the sequence is irrelevant to acceptance or rejection of their successors, then this sequence is being regarded as constituting an argument. (163.)

This seems right: the mark of an argument is that in it a case of some sort is being made, in it some attempt is being made to adduce reasons in favour or against a particular claim. It follows that one can identify an argument by isolating both reasons and the case which is being made

or criticised. That is, one identifies an argument by locating the reasoning which is taking place there, regardless of whether the reasoning is good or bad. Saying this much allows for the possibility of a single deductive or inductive inference counting as an argument. Typically, however, when we speak of deductive or inductive arguments--or when we speak of deductive proofs--we have in mind a string of particular inferences, an extended piece of deductive or inductive reasoning.

Given these characterisations of 'inference', 'logic' and 'argument', what are the salient features of deductive inferences, logics and arguments? First and foremost, deductive reasoning, when it is valid is demonstrative reasoning. An argument is demonstrative just in case the premisses, if accepted to be true, provide complete justification for believing the conclusion to be true as well. Deductive reasoning, as I shall say from time to time, is perfectly conclusive and utterly compelling reasoning. This property of deduction and the crucial notion of deductive validity are closely related. As is well known, there are two standards of deductive validity, or one standard that can be put in two ways: syntactically or semantically. We can say that a deductive argument is valid (in a formal system S) just in case the conclusion is deducible from the premisses (and axioms of S , if any) in virtue of the rules of inference of S ; or, we can say that a deductive argument is valid just in case it is impossible that the premisses should be true, but the conclusion false. The semantic version, moreover, can be expressed in other ways, e.g. a deductive argument is often said to be valid if and only if every model which satisfies the conjunctive set of premisses also satisfies the conclusion. But the upshot of each of these definitions is the same: a deductive argument is valid if the conclusion is demonstrated by the premisses.

The second property of deduction which can be brought out introduces an issue about which I shall have more to say in the next chapter. Following C.S. Peirce (Vol. II, sections 680 and 709), we may say that a valid deduction is explicative, in the sense that the conclusion must be such that any additional information which counts against it must also count against the conjunction of the premisses. Intuitively, this feature of deduction means that the conclusion of a valid deductive argument cannot be stronger, more informative, than the premisses: what the conclusion informs us of, the premisses, at least implicitly, also inform us of. (In the next chapter this feature of deduction will be discussed in terms of J. S. Mill's view of a deductive argument as one the conclusion of which is 'contained in' the premisses; there I will try to show that this way of characterising the explicative feature of deduction is infelicitous.) It is clear that this feature of deduction counts as an explanation for why deduction is demonstrative: it is impossible that the conclusion of a valid deductive argument should be false while the premisses true because the conclusion expresses no new information not already expressed by the premisses.

These two features jointly capture the nature of deductive reasoning: when one validly reasons in a deductive manner one draws out the consequences of a set of premisses in a way which is perfectly conclusive and utterly compelling. There are, of course, other features of arguments, features which even valid deductive arguments must satisfy in order to count as satisfactory arguments--there are, after all, fallacious but valid deductive arguments. But these other features, though important to the judgement of when a piece of reasoning is a piece of good reasoning, are typically not embodied in a characterisation of deduction--a fallacious but valid deductive argument is still a

deductive argument. (I shall be concerned to bring out these other features of arguments in the next couple of chapters.)

Deduction is easily characterisable; induction, on the other hand, poses some special problems. The term 'induction' has been used to designate various rational techniques and argument forms. One thing is clear, however, and that is that inductive arguments are both non-demonstrative and non-explicative, or--borrowing from Peirce again (ibid.)--ampliative. Inductive arguments are non-demonstrative because the premisses confirm or support the conclusion; they make it reasonable to believe the conclusion, they count as reasons in favour of it; but they do not entail it. To argue from instances of a regularity observed in the past to the prediction that that same regularity will be observed in the future, is not to demonstrate that that regularity will be observed. And inductive arguments are ampliative in the sense that the point of an inductive argument is to reach a conclusion which is stronger, contains more information, than all of the premisses taken together. Thus, one of the distinguishing characteristics of inductive reasoning (and non-deductive reasoning generally) is that even with the best inductive argument it is always possible without contradiction to believe that the premisses are true but the conclusion false. In other words, with induction there can be facts and discoveries that can count against the conclusion without thereby counting to an equal degree against the premisses. With induction, although the arguments are not perfectly conclusive and utterly compelling, the conclusion is always intended to advance our knowledge beyond the premisses. Inductive reasoning is reasoning to conclusions which state expected information; induction is a rational form (perhaps the rational form) of expectation reasoning.

But if there has been agreement among philosophers since Bacon that inductive reasoning is both non-demonstrative and ampliative, there has nonetheless been some disagreement as to the form and the range of applicability of inductive reasoning. In particular, two claims about induction form the core of the controversy: (1) induction is always an inference from particular propositions to a generalisation, an inference from a number of propositions having the form 'This P is a Q' to the generalisation 'All P's are Q's; and (2) induction is always an inference from empirical evidence to empirical hypotheses. The first claim can be traced to Aristotle's definition of epagoge as a process of establishing a universal fact, not by deduction from a wider principle, but rather by an appeal to particular cases in which its truth is manifest ((1), Prior Analytics, 68b 15 ff.). The second claim grew out of Bacon's conception of induction as a means for making scientific discoveries and has been carried on into modern times via J. S. Mill's treatment of induction in A System of Logic. It is common to find both claims packed into a definition of induction; hence R. B. Braithwaite defines induction as "the inference of an empirical generalisation from its instances, or of a scientific hypothesis from empirical evidence for it" (257, fn.). Still, increasingly, philosophers in their treatment of induction have felt compelled to deny both of these claims and have argued for a broader notion of induction.

The motivation for this broader notion will be set out in Chapter IV where this conception of induction is advanced and argued for. Suffice it to say at this point that this broader notion retains both of the essential features of induction, and is faithful to the description 'expectation reasoning'. Moreover, this broader conception has the important theoretic advantage of making it possible for us to clearly

distinguish between questions about when propositions provide inductive support for other propositions and questions about the truth or plausibility of the inductive premisses themselves, i.e., questions about the acceptability of the inference. There is, however, a still broader conception of induction which should be mentioned here.

In his Introduction to Logical Theory, P. F. Strawson claims that "reasoning from one non-necessary statement (or conjunction of statements) to another, in which the first does not entail the second...is generally called induction" (237). Given that Strawson presumably does not want to suggest, what he seems to say, that invalid and fallacious deductive arguments are inductive, we might take the liberty to revise his definition in this way: inductive arguments are those which offer evidence in favour of a non-necessary conclusion, which evidence does not entail the conclusion, nor is purported to. It is, it seems to me, unlikely that this is what philosophers generally mean by induction, but it is at any rate a plausible view. The view is simply that induction is non-demonstrative reasoning, that is, non-fallacious, non-deductive reasoning.

It should be clear that I can hardly hold to this conception of induction while at the same time arguing that Deductivism is wrong and that there are rationally legitimate ways of reasoning that are neither deductive nor inductive. Now, oddly enough, although Strawson's characterisation is extremely broad, it is not broad enough to allow as inductive arguments examples of reasoning which are inductive on the moderately broad view of induction I am adopting. For, as I shall argue later, arguments from mathematical induction are inductive arguments, though, having both necessary premisses and a necessary conclusion, they fail to count as inductive on Strawson's view. In any event, I shall

make an attempt later to distinguish induction from other forms of non-demonstrative reasoning, though the manner in which I shall make the distinction will be such as to make it very unlikely that a hard and fast line of demarcation can be discovered, or if arbitrarily drawn, justified.

These, then, are the general characteristics of deduction and induction. Throughout this thesis my perspective on deduction and induction will be philosophical. I shall not be offering systems of deductive or inductive reasoning. I am taking these logics as given, familiar and, indeed, unproblematic. I have no intention to launch an attack on logic. The sceptical pose I adopt in the next chapter, for example, is not designed to cast doubt on the legitimacy of deduction as a form of reasoning; rather, it is an attempt to discover why deduction is a legitimate form of reasoning, and, more importantly, how we might go about showing that it is. Nor am I going to propose some sort of 'alternative logic'--a bizarre or deviant logic which can be produced by tinkering with the standard systems. My concern here is with what it means to reason, and specifically, with what it means to reason about questions which do not seem to be answerable either deductively or inductively.

I have claimed that Deductivism is imposed on philosophical problem situations; but I have yet to give an example of what I mean by this 'imposition'. What I want to offer, then, is a case in point. I want to consider the effect on the question of the nature of moral reasoning which Deductivism has had in recent ethical theory. Deductivism is, I suggest, the received view, the standard treatment of reasoning; and when it is brought to the issue of moral reasoning it tends to turn the discussion towards a debate over the justifiability of moral

principles. To a surprising extent, meta-ethical positions which represent the spectrum from sceptical accounts to objectivistic accounts share the view that, if it is possible to reason about moral questions at all, the reasoning which takes place requires moral principles. And this common bias is the product of Deductivism.

Deductivism in Ethics

I want to begin with a sort of composite picture of the standard treatment of how one is to proceed when one deliberates over, or argues about a substantive moral question. (The details and loose ends of this account can be omitted or left to dangle in order to make the rough outlines of the received view stand out a bit more clearly.) What we are supposed to have in the standard case of a moral conflict is two or more moral judgements which are at odds, e.g., 'You ought to pay your taxes' and 'It is my moral duty not to pay my taxes'. The conflict--and let us suppose it takes the form of a verbal argument--is then seen to continue, step by step, with each side attempting to justify, or prove, their original judgements. On the received view the pattern of justification is straightforwardly deductive, or possibly, as R. M. Hare has suggested, hypothetico-deductive (cf. Hare (2), 89-90). That is, each party to the dispute offers up a logically more inclusive (or, deductively stronger) judgement which, if acceptable, would justify the original judgement by virtue of the fact that the original judgement is entailed by it. So: 'You ought to pay your taxes because you are a citizen of this country and every citizen ought to pay his taxes'; or 'It is my moral duty not to pay my taxes because 90% of the money goes to finance a genocidal war and everyone has a moral obligation to oppose in any way they can such a war'. The logic of such reasoning is very simple indeed,

taking the form of elementary deductive inferences, such as universal instantiation, modus ponens, and a handful of others.

It is easy enough to see that the general direction which this sort of deductive justification takes is from the particular to the universal. Justification for a judgement about what some moral agent ought or ought not to do in this case is sought from other judgements about what everyone ought or ought not to do in every case like this one. These general judgements go by various names--'moral principles', 'moral rules' and 'moral standards' being the most common. There are, moreover, degrees of generality in evidence here, so that justification for a moral principle may be forthcoming from another moral principle which is more general still (a moral principle, that is, which is deductively stronger than the first). Moral reasoning, then, is seen as reasoning which depends on an appeal to principles, rules or standards.

What purposes are served by an appeal to a principle in the context of a moral argument? At least three purposes are generally recognised. The first, and most straightforward purpose is that of providing reasons for a moral assertion. If a moral principle is accepted by both parties to the dispute, and if one of the original judgements is entailed by it, then this fact provides evidence for that judgement. Thus, 'Every citizen ought to pay his taxes', a moral principle of a fairly low level of generality, counts on the received view as a reason in favour of the judgement that you or I ought to pay our taxes, given that we are citizens, given, that is, that our case falls under the principle. So, moral reasoning which follows the deductive pattern of appealing to principles is viewed as a rational procedure of reason-giving.

A second purpose has to do with other consequences of the moral principle to which one is appealing. Specifically, being general in scope,

the moral principle entails that other moral agents satisfying the antecedent conditions have the same responsibilities as you or I. If we bring up these other cases, then the full significance of the particular judgement in question can be revealed. If I object that I ought not to pay my taxes because I am poor, then one counter-reply might be that other poor citizens have this duty so my poverty is not a reason for an exemption in my case. On the other hand, my objection can be supported by altering the antecedent conditions of the principle, i.e., rejecting 'Every citizen...' in favour of 'Every citizen, who can afford to,...'. This move has the effect of altering the significance of my original judgement, and so, of changing, however, subtly, the nature of the dispute.

The last purpose served by principle citation is similar. By citing a principle one is reminding one's opponent of the requirement of consistency. One can be faulted for accepting a universal principle while at the same time refusing to acknowledge the application of it to the case in question. I cannot agree that every citizen ought to pay his taxes, claim that therefore you must pay your taxes, while refusing to accept that obligation in my own case (if I admit that I too am a citizen). The argument here can be approached from the opposite direction, and when it is it becomes an argument based on what Hare has called the 'Universalisability Requirement' (*Ibid.*, e.g. 10ff.). If I make a particular moral judgement which makes an essential reference to some agent or act, then I must also accept the universalisation of that judgement (universalising over all moral agents and all acts of that type). As a consequence, I am forced to accept the universalised judgement, the moral principle. If I claim that you ought to pay your taxes because you are a citizen (i.e., if that is the reason that I am advancing for your obligation), then I am implicitly committing myself to the claim that all citizens have that

duty, and hence, if I am a citizen, that I do as well. Maintaining that I do not have the duty, while acknowledging my commitment to the moral principle, opens me up to the charge of inconsistency.

This, then, is the composite picture. One of the more conspicuous features of it is that reasoning about moral questions tends to push one in the direction of ultimately general, indeed axiom-like, moral principles. These principles, by their very logical nature, cannot be justified in the way in which they justify their deductive consequences, since they are, as it were, as far 'down' the deductive path as one can go. Moral reasoning presupposes a bottom, foundational level of moral principles--this is a built-in feature of the Deductivist approach. It is easy enough to discern this particular thesis about moral reasoning in most non-cognitivist and subjectivist accounts of morality.

Consider the popular view that moral conflicts have a composite nature, being composed of conflicts over facts and conflicts over values. Conflicts over facts, the story goes, can be resolved easily enough, we merely have to agree on what the facts of the case are. But after agreement at this level has been reached, any further reasoning is thought to require an appeal to a common set of general moral principles. A. J. Ayer, for example, writes that "argument is possible on moral questions only if some system of values is presupposed" (111). A system of values is a set of moral axioms which can be brought to an argument. Failure to pay one's taxes if one is a citizen is wrong because it is a moral rule (say) that every citizen ought to pay his taxes. Nothing could be, deductively, simpler. But if the rule that every citizen is so obliged is questioned then one must search for another principle which will have this rule as its consequence. With a little ingenuity one might be able to show how this rule is a deductive consequence of a principle of great

generality, say, 'One ought to promote the welfare of others'. But then one is at, or very near, the level of one's moral axioms, one's system of values. If there is disagreement at that level, argument grinds to a halt: "...we do not and cannot argue about...the validity of these moral principles" (Ibid., 112).

But if we cannot argue about these principles, how is it that they play any role in justification at all? At this point the basic subjectivist doctrine steps in to fill in the gaps between the Deductivist account of moral reasoning and the surrounding ethical theory. These principles are the expression of moral preference, they represent our commitment to ultimate moral beliefs:

It is a commonplace of moral philosophy that at the highest level of justification for any given context, at the level...where there are no further reasons to be given for the holding of a principle, one just holds the principle or maintains the purpose, one just approves or has what has been called a pro-attitude. (Montefiore, 143.)

Making such a fundamental moral judgement, it is thought, is exercising one's freedom as a moral agent to live by the principles of one's choice. In addition, one is also setting up the basis for deductive justificatory moves. Moreover, as long as one is consistent in the principles and other moral beliefs one holds, any moral judgement can be deductively justified if one makes the commitment to the appropriate principle.

But we can and do question moral principles, fundamental or not. We do, but, according to the subjectivist, once the facts are agreed on and the deductive relations laid out any further discussion must be a form of non-rational persuasion. If a moral conflict is not rooted in a disagreement over factual belief, then it is a disagreement of attitude and no reasoned solution to the conflict is possible. For central to the subjectivist view is the doctrine that deductive inference can only carry one from empirical claims to empirical claims, or from moral

judgement to moral judgement. That one cannot deduce a moral judgement from empirical facts is the basis of the so-called 'is-ought' problem: There is thought to be a gap between statements of fact and statements of value or obligation, a gap the bridging of which constitutes a logical error. Needless to say, it is this gap which makes it necessary for the moral reasoner to move back towards his basic moral principles. Radical disagreements over moral questions--moral conflicts which persist--are accounted for by the existence of different moral axioms and different systems of values.

The subjectivist is thus committed to the view that rational discussion of moral issues is possible up to, but not including, ultimate moral principles. As far as the Deductivist model of justification is concerned, of course, it does not matter where these basic moral principles come from, as long as they are available for use. Some philosophers, however, while agreeing on little else, feel that there is something very wrong with the view that one can choose one's moral principles. It has been suggested, for example, that morality must be seen as a collection of moral systems or practices which individual moral agents participate in:

In order to make a moral judgement one must belong or be related to a moral practice within which, quite independently of any decision on the part of those who belong to that practice, certain facts entail that some things are right and wrong. (Phillips and Mounce, 14.)

The governing idea behind this relativistic view is that the very meanings of moral terms, the criteria for the application of moral concepts, are given by the moral practice one has been brought up in. One cannot decide to call lying or murder wrong--if one has been brought up in, say, the Judeo-Christian tradition--since lying and murder are the sorts of acts that those of us who participate in this moral practice apply the word 'wrong' to. It is correct to say, that is, that our moral

practice determines how we see the acts we perform; it determines as well what counts as a moral issue and what sorts of considerations count as moral reasons (vide, Beardsmore, 24 and 34). Since our moral practice provides us with the criteria for applying moral concepts, there will be within the moral practice entailments between facts and moral judgements. The gap between the facts of a particular situation and the judgement that what happened was a case of, say, murder, as well as the gap between that essentially descriptive judgement and the further, evaluative, judgement that it was the morally wrong thing to do, both of these gaps are legitimately bridged for those of us sharing a moral practice. These gaps are bridged by the very meaning, for us, of the word 'murder'. Yet:

To maintain that within a moral practice certain facts will entail certain moral conclusions does not preclude the possibility of there being different moral practices within which the same facts entail different conclusions. (Phillips and Mounce, 15.)

The deductive relations obtain, but only relative to a particular moral practice. Moral judgements can only be understood--the moral terms employed in them can only have meaning--as they appear within moral practices. These practices will differ under different social conditions, so there is no a priori limit to the number of different forms moral judgements can take (Ibid., 61).

As far as the logic of moral reasoning is concerned, this view has the consequence that there are two quite distinct forms that moral conflicts can take. One sort of moral conflict takes place within a moral practice. For such a conflict both disputants share what Beardsmore calls a "framework of agreement" by virtue of the moral practice they share. This framework of agreement is built out of a shared understanding of moral concepts, and is manifested in an agreement over moral principles.

Since there is agreement at this level, the disagreement which results in conflict must be a disagreement over the non-moral facts, or, alternatively, the consequences of the shared moral principles. Such moral conflicts can be resolved, or at least, there is a solution in the offing. But these sorts of moral conflicts must, on this view, be sharply distinguished from moral conflicts which arise between moral practices. When there is no single moral practice in the background of a moral conflict, then the framework of agreement is absent. In such cases, the arguments are not decidable, even in principle "simply because the disputants cannot even agree over what criteria to apply. They each have reasons for the judgements that they make, but neither admits the relevance of the other's reasons. The argument has reached deadlock" (Beardsmore, x). Reasoning is crippled for such cases since deductive procedures cannot lead to a single, unambiguous conclusion if the meanings of the predicates employed in the propositions serving as premisses are not held constant. But the crucial moral predicates derive their meaning from moral practices, and when two disputants participate in two different moral practices these predicates acquire different meanings. We cannot, by employing deductive procedures, adjudicate between moral practices.

The views so far mentioned all incorporate Deductivism, and the result of this is the same in each case: the limitations which these views have placed on moral reasoning are just the built-in limitations of deductive reasoning; namely, the limitations which result from the need for an axiom-like foundation and for unambiguous predicates. The need is felt for major premisses, premisses from which particular moral judgements can be deduced and hence to which we can appeal in order to justify those particular moral judgements. Alternatively, the need is

felt to bridge the supposed gap between questions of fact and questions of value and obligation, to find some way to deduce moral conclusions from non-moral evidence.

Now these needs are felt in ethical theory, I want to contend, primarily because the deductive model of reasoning has been assumed to capture moral reasoning. We have looked so far at some common meta-ethical positions, non-cognitivism, subjectivism and a recent form of relativism; but the assumption of Deductivism can be seen on the other side of the meta-ethical fence as well. With moral objectivism, moreover, a wider range of theoretical consequences of Deductivism can be discern. In Kantian ethical theory, to take an obvious example, the standard Deductivist appeal to principles takes on a two-fold nature. Firstly, as moral agents deliberating a course of action we will, if we are rational and if our acts are purposeful, act in accordance with a maxim, a personal rule of conduct. Secondly, we will also appeal to the a priori Categorical Imperative: "Act only on that maxim through which you can at the same time will that it should become a universal law" ((1), e.g., 21). Since the Imperative does not refer explicitly to kinds of acts (and thus is, in Kant's terminology, 'pure'), it cannot serve as a major premiss in a moral argument unless it is mediated by a maxim. But the crucial feature of the Imperative is that it is foundational, the essential ingredient of the 'meta-physic of morals'.

Although I cannot hope to be able to cover all possible objectivist views here, two other objectivist programmes should be mentioned: naturalism and intuitionism. Each in its own way attempts to move away from a strictly deductive model of moral reasoning; yet, Deductivism is still very much in the background.

Naturalism in ethics, very roughly speaking, is the view that the key moral concepts can be explained, even defined in empirical, scientifically acceptable terms. In its extreme, reductive form, naturalism suggests that the goal of ethical theory should be to expose as clearly as possible those actual human values about which all humanity is in agreement. This suggestion leads quite naturally to a general programme for delimiting the range of values that there are, a programme which relies on inductive procedures. This programme has been vigorously defended by Abraham Edel who has repeatedly argued that ethical theory should be patterned after a scientific theory (vide, e.g., (1) 74 and 111; (2) 131 f.). Although Edel's programme constitutes an extreme form of ethical naturalism, it shares with most forms of that view the assumption that subjectivism and relativism in ethics are false because, among other reasons, there are ways of bridging the supposed logical gap between statements of fact and judgements of value and obligation. The goal of Edel's programme is to establish law-like generalisations in ethics so that these may serve as major premisses in essentially deductive ethical arguments. The assumption of Deductivism is operating in full force here: reasoning from principles is the undisputed basic mould of reasoning into which moral argumentation is to fit. Edel's naturalism is designed to provide suitable major premisses for deductive arguments.

The second objectivist programme worth mentioning here brings out another philosophical variation on a common theme, another plausible response to the consequences of Deductivism. Ethical intuitionism is the view that at some point in the justification of a moral assertion one must rely on a kind of direct, non-inferential cognition which is not based, wholly or in part, on reasoning. The link with Deductivism can be seen in the works of most intuitionists. One may cite as an example

H. A. Prichard's 1914 paper "Does Moral Philosophy Rest on a Mistake?". The mistake which Prichard argues moral philosophy does indeed rest on is the suggestion that it is possible to prove, or to give good reasons for, claims about moral obligations. We cannot, Prichard argues, acquire a knowledge of right and wrong by reasoning things out; rather the rightness or wrongness of a particular action "arises in our unreflective consciousness" (16). Prichard's argument relies on the assumption that if there were reasoning in ethics, it would be deductive in form. The plausibility of Prichard's case depends on the familiar claim that it is impossible to deduce obligations to do things in particular cases from other, non-moral statements. The resort to intuition is thus an attempt to underwrite the objectivity of morals in the light of what is taken to be the ineffectualness of reasoning.

Generally speaking, the need for intuition has been thought to follow from the fact that at some point in the justification of a moral assertion one must be able to appeal to justificatory solid ground. If that solid ground cannot be axiom-like moral principles, then it must take some other, ultimately secure form. For the intuitionist, moral intuitions serve to bring the request for justification to an end. The assumption here is that only justification which comes to a definite end, in, moreover, a finite number of steps, is genuine. That assumption is just a consequence of the Deductivist model of justification as an appeal to general principles.

But it is important to emphasise that the intuitionist is forced to claim both that moral intuitions are self-evident and that they cannot be proved. And that they are forced to make both claims puts them in a very awkward position. One manifestation of the tension between these two claims is particularly characteristic of modern intuitionism. One

senses a kind of faith that the truth will eventually win out. At one point in his The Right and the Good, W. D. Ross expresses the intuitionist's faith in these words:

To me it seems as self-evident as anything could be, that to make a promise, for instance, is to create a moral claim on us in someone else. Many readers will perhaps say that they do not know this to be true. If so, I certainly cannot prove it to them; I can only ask them to reflect again, in the hope that they will ultimately agree that they also know it to be true. (20, fn. 1.).

But one really cannot help but think something is very odd about this. If it is self-evident that making promises creates a moral claim on us, then one would expect that that would be very easy to prove, that it would be possible to prove, or, at the very least, that it would be possible to give good reasons for thinking it so. The oddity here is a consequence both of the influence of Deductivism and an overall pessimism about the effectiveness of reasoning--a pessimism which is a consequence of Deductivism. Because intuitionism does not abandon the Deductivist model of justification and proof, it requires an ultimate ground of justification which is itself beyond justification. To provide this, moral intuition must be able to yield self-evident moral assertions of one sort or another, either assertions about particular obligations or, as Ross thought, about prima facie obligations. But, since reasoning must end at the foundations on the Deductivist model, no further reasoning is possible at this level and these same self-evident moral assertions cannot themselves be proved. One must, it would appear, resign oneself to the fact that the scope of reasoning cannot extend to the foundations themselves.

I have not, of course, begun to do justice to the many variations of the basic meta-ethical positions I have touched on above. What I hope to have shown, however, can be shown without going into too much detail,

and without trying to cover every variation and refinement. In brief, the assumption of Deductivism tends to bias the discussion of the nature of moral reasoning in certain conspicuous ways: Justification of moral assertions is seen to be a matter of appealing to moral principles. And the justification of ultimately general principles, no longer a deductive matter, seems to require either induction or intuition. At the level of ultimate principles, moreover, some philosophers resort to the view that it is merely a matter of personal choice what moral principles there are. Others suggest that some of them are analytic, others that ultimate moral principles are ultimate constituents of moral practices to which people belong or adhere. What we seem to have in contemporary ethical theory is an across the board insistence that reasoning in morality must be reasoning from principles. I want now to consider the nature of reasoning from principles.

CHAPTER II

REASONING FROM PRINCIPLES

In what follows, I shall be endeavouring to set out and develop several considerations in favour of the rationality and logical priority of reasoning which is neither deductive nor inductive. The keynote of this discussion is that non-deductive and non-inductive reasoning is ordinary reasoning--it is not so much a method of arguing as what we do when we reflect and deliberate; it is not some exotic, alternative logic, it is being rational. The difficulty will not be in finding cases on non-deductive, non-inductive reflection and argument, but in appreciating situations in which we unself-consciously argue and deliberate in a non-deductive, non-inductive reflection and argument, but in appreciating situations in which we unself-consciously argue and deliberate in a non-deductive, non-inductive fashion whose familiarity impedes recognition. The thesis is not being advanced that deductive or inductive reasoning--either in the context of formal, logical calculi or in the context of informal reasoning--ought to be supplanted by some third sort of reasoning. Rather, I want to argue that deduction and induction derive their rational strength--their usefulness and their legitimacy--from a logically more fundamental manifestation of proceeding in a rational manner which itself does not exemplify deductive or inductive patterns. The value of deductive and inductive reasoning in formal contexts will be argued to rely on certain features of the conceptual makeup of areas of human knowledge which are amenable to a degree of systematisation.

In the present chapter, I want to restrict the discussion to deductive reasoning, and in particular to the model of deductive reasoning which the Deductivist favours, namely the model of deductive justification from principles. The discussion will take the form of two formulations of a dilemma for deduction as a form of reasoning. The first of these, which employs J. S. Mill's famous contention that deductive reasoning always involves a petitio principii, will be shown to fail to pose a genuine problem for the status of deduction as a form of reasoning. Yet, the insights from Mill's proposed solution to the difficulty he thought he saw in deductive reasoning will motivate a second formulation of the dilemma which does not fail. I shall then attempt to resolve the dilemma, and in the process reveal a crucial feature of reasoning, that reasoning rests on the recognition of the legitimacy of particular and informal examples of reasoning. It will then be possible to isolate those features of the structure of formalisable subject matters upon which the legitimacy and value of formal deductive reasoning from principles rests. The dilemma, and its resolution, will yield for us the following two results: in reasoning, the informal is prior to the formal, and the particular gives life to the general.

A Dilemma for Deduction--First Formulation

We require of a piece of reasoning both utility and legitimacy: we want not only an epistemic advance of some sort but also an assurance that that advance is rooted in what we already know or believe. These requirements are intertwined in at least two classical definitions of reasoning. Aristotle defines reasoning as "an argument in which, certain things being laid down, something other than these necessarily comes about through them" ((1), 100a 25-26). For Aristotle legitimacy is

assured by a necessary connection which links the result of the reasoning with what has already been laid down and is not in question. But there is an advance since it is "something other" than what has been laid down which comes out of the reasoning. And, speaking more generally, Cicero defines an argument as "a course of reasoning which firmly establishes a matter about which there is some doubt" (Book I, section 8). For Cicero the advance is a removal of doubt, and the legitimacy of the reasoning is secured by the matter being firmly established.¹ But consider as well Charles Sanders Peirce's claim that the two aims of the logician are "1st, to bring out the amount and kind of security (approach to certainty) of each kind of reasoning, and 2nd, to bring out the possible and esperable uberty, or value in productiveness, of each kind" (Vol. VIII, section 384). Security and hoped for uberty, legitimacy and epistemic advance--both epistemological requirements are clearly part of our conception of what it means to reason. Without epistemic advance, removal of doubt or a new understanding, the legitimacy of a piece of reasoning does not suffice to make it worthwhile; and without legitimacy we have no reason for viewing what has been produced as an addition to our knowledge or a new understanding.

But then, how can a deductive argument qualify as reasoning, since legitimacy (i.e. formal validity) rests on a feature of deductive inferences which seems straightforwardly to preclude the possibility of epistemic advance: whatever is required to see how the premisses can be true together suffices to show that the conclusion is also true. In a deductive argument, as it is often put, the premisses contain the conclusion. But if a deductive inference is valid, we have not found out something other than what we knew before; and if we have learned

something from the conclusion which was not already there in the premisses, then the inference was not valid. How can utility and legitimacy be reconciled in the case of deductive inferences?

The key to this dilemma--if there is one here--obviously lies with the observation that the conclusion of a deductive inference is contained in the premisses. Yet there are various sorts of deductive inferences, which go to make up deductive arguments of different degrees of complexity, and the complexion of this dilemma changes with the example. If we view deduction as reasoning (genuine or sham) from the general to the particular and take as our paradigm of a deductive argument a syllogism of the following sort

(1) All men are mortal.

 Socrates is a man.

 Therefore: Socrates is mortal.

then we might be inclined to agree with Mill that every deductive argument is a petitio principii (II, 3, 2): The particular case of Socrates, Mill argued, is either included in the general major premiss or it is not. If not, then it is not legitimate to move to the conclusion; but if it is (and this surely is what it means to say that all men are mortal), then whatever assurance we have that the major premiss is true must rest on a consideration of all cases of mortality, including, of course, Socrates's. Hence, the conclusion of (1) does not appear to be an epistemic advance at all. Suppose we take Mill's position towards deduction to be the basis for the first formulation of a dilemma for deduction. Does Mill have a legitimate complaint?

It is important, first off, to see how much of Mill's challenge rests on his choice of examples, namely syllogisms like (1). For

suppose we change the example. And for a good contrast, suppose we move from the syllogism to an argument consisting of a series of inferences which is part of an axiomatised system of mathematics. Consider the deductive argument whose conclusion is

(2) Every Cauchy sequence of rational numbers converges.

and whose premisses are the axioms of Zermelo-Fraenkel set theory and a handful of definitions introducing notational conventions. For this example it surely would be false to say that the conclusion, (2), doesn't tell us something new, something different from what we already know by virtue of knowing the axioms and definitions. Yet, at the same time, (2) is 'contained in' the axioms and definitions in the appropriate sense:

(2) is a legitimate deductive consequence of these axioms and definitions, and nothing else, i.e. there is a valid demonstration of (2). Still, it would be at least odd to call this deductive argument a petitio.

But why would it be odd to call it a petitio? Is it because we can not detect in the axioms and definitions the presence of the theorem about Cauchy sequences? This would appear to be something which we can do in the case of (1). But in (1) sense can be given to talk of detecting the conclusion in the premisses. When there is a move--or a series of interlocking moves--from axioms and definitions to (2) it is not clear what it is that we are supposed to detect. Rather it would seem that here we have a real advancement in knowledge--we might have had the suspicion that every Cauchy sequence converges (since they are designed with that idea in mind), but now we know. Indeed, any complex deduction like this one shows up part of the artificiality of the dilemma as I have stated it. We can only be said to know the deductive consequences of the statements we already know if we know that they are deductive

consequences. Sometimes there are enormous gaps between what we know and what deductively follows. Yet it is important to note that the necessary connection which links the Cauchy sequence theorem with the conjunction of axioms and definitions--as these statements are embedded in a suitable axiomatic theory--is mediated by a very large number of separate steps, each of which is a sub-argument taking as premisses the conclusions of preceding sub-arguments (or else the axioms themselves). The reason why there doesn't appear to be a petitio here is that the 'distance' between the premisses and the conclusion is too great. Of course that large gap can be filled by countless sub-arguments--most of which having the same degree of deductive simplicity as our first example, (1). Yet, even if we were to actually travel that chain of separate inferences, we would doubtless be amazed at the great epistemic distance there is between the axioms and the theorem.

It is legitimate to object at this point that to say that there doesn't appear to be any question begging going on here doesn't really discredit the claim that there really is. But there is another consideration which makes Mill's general claim doubtful. There is a paradox which arises when one attempts to argue Mill's case for examples of complex deductions (especially complex deductions taking place in well-defined formal systems like first-order functional calculus); and the paradox is this: we have no way of determining in a non-question begging way whether these deductions are or are not themselves questions begging. This is because the only way in which it could be shown that the truth of a putative theorem was being begged (and the only way in which it could be shown that the conclusion is in the premisses) is precisely the way it could be shown that the putative theorem is a theorem. But

then, how do we know that what we are showing is the begging of a question or the theoremhood of an expression? In short, when the distances are great and the deductions complex we have no independent means of making sense of Mill's accusation that every deductive argument is a petitio.

It might seem unfair to bring up examples like (2) to discredit Mill's claim. He was concerned with deductive arguments like (1), not with very complicated arguments taking place in formal theories, the existence of which he was, for the most part, unaware. Moreover, as we just saw, complex arguments resolve themselves into series of simple sub-arguments like (1). So, ultimately, it is to these simple moves that the question of petitio must be directed.

But even at this level Mill hasn't got it right. The example Mill chose is particularly suited to the point he wanted to make, since in asserting the major premiss that all men are mortal we are also asserting that Socrates is. Yet, other examples of deductive moves which are on a par with (1) are not as easily handled in the same way. Consider the following:

(3) Jones is either a saint or a sinner.

 Jones is no saint.

Therefore: Jones is a sinner.

(4) Jones is a sinner.

Therefore: There is a sinner.

In the case of (4) the conclusion makes no reference to Jones, the premiss does not assert the conclusion. And for (3) we would be making an adventurous claim indeed if we suggested that Jones's status as a sinner is being begged by the major premiss. The possibility that Jones is a sinner can indeed be detected in the major premiss. But, to say that

Jones is either a saint or a sinner is neither to assert that Jones is a sinner nor is it to assert that he is a saint. The conclusions of (3) and (4) are not--what we might call--assertorical parts of the respective sets of premisses. Moreover, if the claim that question begging is present in these examples were to be pushed, we would soon run up against the same paradox as before: to show that the conclusions of (3) and (4) are built into the two sets of premisses could only be shown by showing that the two conclusions follow deductively from the premisses. If we had other reasons for being confident than showing that a conclusion follows deductively from a set of premisses is showing that there is question begging going on, then we could agree with Mill. But without these reasons, we would be begging the question of whether (3) or (4) is a petitio if we viewed deductive inference as a sign of question begging.

The dilemma we posed for deductive argument as genuine reasoning can not rest on the rash claim that every deductive argument is a petitio principii. We have allowed Mill a case for deductive inferences like (1) because some sense can be given to a non-question begging detection of the inclusion of the conclusion in the premisses--namely that the conclusion is an assertorical part of the premisses. But the simple picture of deduction as inference from the general to the particular--a picture which Mill never questions--is in fact very misleading. Only in cases where an instance of a universally quantified premiss is--or immediately leads to--the conclusion can real sense be made of the claim that the conclusion is implicitly or surreptitiously being asserted by the premisses. Universal instantiation is, of course, one pattern of deductive inference; but it isn't the only pattern. There is, in deductive inference, always

a move from relative deductive strength to less or equal deductive strength, but that move is not always the same as the move from general to particular. Specifically, the general-to-particular picture misleads by suggesting that in knowing what the general premisses mean we have encountered and acquainted ourselves with the particular conclusion. There can be real epistemic advance which results from a valid deductive argument. But, Mill thought that as well, and it is important to see how he argued for real advance given the sort of example he was considering.

Mill fairly clearly saw the seeds of the dilemma for deduction. He put it this way. How do we reconcile the fact that "from a general principle we cannot infer any particular, but those which the principle itself assumes as known" with the fact that "a large portion of our knowledge" is acquired by means of the syllogism (II, 3, 2)? We can arrive at the knowledge of the as yet unobserved truth that Richard Nixon is mortal by means of a syllogism of the same form as that used to arrive at the already observed and recorded fact that Socrates is mortal. The claim that R. N. is mortal is an epistemic advance; but how could it be if it has already been asserted by the major premiss? Mill's answer is that there is, in fact, no inference from 'All men are mortal' to 'R. N. is mortal', rather the inference has already been made once 'All men are mortal' is asserted. The universal proposition is both a record of particular instances of mortality which have been observed in the past, and an inductive inference to instances of mortality which have as yet been unobserved (as well as a retrodictive inference to past, unrecorded and presently unobservable cases). The particular case of R. N.'s mortality is indeed included in the general major premiss; but there is epistemic advance here because of the inductive move from

particular cases of mortality to the particular case of R. N.'s mortality.

There is an initial difficulty with this solution. And that is that it turns the syllogism into a silly flourish. The syllogism--to use Mill's own metaphor--seems to do nothing more than take us up a hill and down again when we could have stayed in the valley. Mill was obviously very conscious of the objection that his view seems to strip from deductive inferences any function, any real value. To this Mill has a reply.

Mill acknowledges the view that the function and value of deduction is the prevention of inconsistency: the syllogism prevents us from assenting to something the truth of which contradicts something else to which we have already, on good grounds, given our assent. Yet, Mill contends that only his view of the real nature of the inference taking place in a syllogism can make sense of this function. The real reason why we believe that R. N. will some day make his mortality manifest is that R. N.'s father, and his father and many others besides have died. But, the claim that R. N. will not die is not inconsistent with these facts, so Mill must provide some assertion which is inconsistent with the denial of the conclusion of the syllogism. This assertion must have been, Mill insists, the result of induction.

If we suppose that, prior to considering this syllogism, we had reasoned inductively from past cases of human mortality to the present case of R. N., then upon considering this syllogism we would indeed have a reasoned assertion on hand which is inconsistent with the denial of the conclusion. But if we had actually made the prior inductive inference (and remembered that we had), the syllogism would be quite superfluous. Nonetheless, if we had not made the inductive inference to R. N.'s case,

as seems likely, still it would have been proper to do so. As a consequence, the major premiss is just the "memorandum" we make of all the actually observed cases and all the inferences which could have been properly drawn to future (and past) cases. The function and value of the syllogism can thus be seen in our use of general propositions: we employ these propositions to remind us that we should refrain from drawing inferences inconsistent with those which we could have previously drawn. The rules of the syllogism are thus directions for interpreting these general propositions with the goal of maintaining consistency between the present (deductively derived) conclusion and the previous (inductively derived) conclusion. Reasoning itself is fundamentally from particular to particular in an inductive fashion; indeed, we could always reason in that way. Still, there is great value in the use of general propositions since they are, in effect, the record of previous observations and inferences: "Though not necessary to reasoning, general propositions are necessary to any considerable progress in reasoning" (II, 3, 5).

In his attempt to get around the dilemma for deduction as he saw it, Mill has contributed something correct and something confused. In his favour he has brought attention to the ultimate particularity of reasoning: ultimately, all reasoning is from the particular to the particular. The faith we have in general propositions is just the faith we have in the particular instances which are collected together in general propositions. I want to return to this lesson later, where the impact of this claim will be explored. But for the moment it is important to see where and how Mill went wrong.

Mill went wrong in supposing both that deduction itself cannot provide epistemic advance, and that the epistemic advance which seems to

result from deduction actually results from prior, inductive moves. This latter error, as has been repeatedly pointed out, leads to the thesis that all universal propositions--even those of logic and mathematics--are empirical generalisations. This contention can be easily shown to be false: There are an infinite number of deductive consequences of the axioms of first-order functional logic, and there are an infinite number of these which are universally quantified. It is impossible that all of these logical truths could have been the product of actual inductive generalisation. Let A be one which hasn't been the result of induction. Now A is, ex hypothesi, provable and has never been the result of inductive reasoning. So, even if the axioms of this (or any) logical calculus are inductive generalisations, not every universal consequence of them could be. (The same argument could obviously be repeated using axioms of mathematics and a, previously unconsidered, universal consequence of them.)

But what about Mill's other claim that there is no epistemic advance which directly results from deductive inference? This claim is the root of the Millian formulation of the dilemma for deduction. I have already noted that there are exceptions to Mill's claim that every deductive argument involves a petitio, and it was this claim which motivated Mill to look elsewhere for the basis of the utility of deduction. To be sure, as arguments, some deductive inferences would be blatant examples of the petitio fallacy. It is, after all, deductively quite all right to conclude that Jones is a sinner from the premiss that Jones is a sinner. But, the rejection of the autonomy and utility of deduction which is based on the grounds that conclusions are implicitly asserted by premisses cannot be sustained. Moreover, this way of looking at

deductive arguments blinds us to the fact that genuine epistemic advance is possible with deduction. Deductive reasoning is explicative reasoning, to be sure, but frequently the aim of reasoning is to determine what follows from something else, that is, to determine what the premisses of an argument commit us to. Deductive reasoning is not only reasoning which results in proof, it is also reasoning which can result in a discovery. Deductive reasoning does not cease to be explicative--the conclusion is still 'contained in' the premisses--when it is reasoning which results in a discovery about what the premisses imply. Deduction may just be 'putting two and two together', but sometimes the conclusion constitutes genuine epistemic advance.

To sum up, then, I began with the claim that we require of a piece of reasoning both utility and legitimacy. A piece of reasoning is useful (in this context) if it yeilds some sort of epistemic advance, if it brings forward new information, if it justifies, explains, removes doubts or answers questions. The legitimacy of a piece of reasoning, on the other hand, is an issue which brings up the question of whether we are entitled to the epistemic advance. This second requirement demands that the reasoning be defensible, that we have good reasons for allowing the argumentative move from premisses to conclusion. As I originally set it out, the dilemma for deduction went like this: legitimacy, i.e., formal validity, requires that the conclusion of a deductive argument be 'contained in' the premisses; but (i) if a deductive argument is valid, then there must be no genuine epistemic advance possible, and (ii) if there is epistemic advance, then the argument must not be valid, the advance must be traceable to some other source. The upshot of the dilemma is that the two requirements of utility and legitimacy are

are incompatible in the case of deductive inference. Mill's formulation of the dilemma fails primarily because he was caught up in a distorted picture of what accounts for the legitimacy of deductive inference--the metaphor of conclusions being contained in premisses leads, or rather misleads, Mill into thinking that no genuine epistemic advance can result from a deductive inference. But sometimes there is genuine epistemic advance which results from valid deductive arguments which are not in any recognisable way supported by, or consequent on, inductive inferences.

As it happens, the best examples of deductive arguments which produce new information--and often epistemic advance of the most ingenious sort--are those deductive arguments couched in the language of formalised systems of mathematics and logic. The priority of informal deduction notwithstanding, it is within the boundaries of these formalised systems that deductive reasoning manifests itself as an extremely powerful mode of discovery. Indeed, the power of deduction in such contexts has humbled philosophers to the extent that deduction is generally viewed to be the only truly secure and trustworthy form of reasoning. The result has been, of course, dogmas like Deductivism. In short, no discussion of deduction would be accurate or complete without a consideration of formal deductive logic. And this we shall immediately do with an eye to the problem of reconciling utility and legitimacy. The discussion will reveal a formulation of the dilemma in terms of formal deductive reasoning which is, it seems to me, perplexing. This second formulation will account for the initial plausibility of Mill's attack on deduction which was founded, as I tried to show above, on the general-to-particular model of deductive inference.

A Dilemma for Deduction--Second Formulation

The first formulation of the dilemma for deduction turned on the requirement of epistemic advance. Mill thought he had captured the crucial condition for deductive validity when he repeated Bishop Whately's catch-phrase: 'the conclusion is contained in the premisses of a valid argument'. Mill went on to account for the possibility of epistemic advance in the light of this conception of deductive legitimacy. Here I want to formulate a genuine dilemma for deduction which turns on the requirement of legitimacy.

As was noted in the last section, this second requirement of reasoning demands that reasoning be defensible. Now a deductive argument is legitimate when it is valid, and deductive validity rests on a very good reason for allowing a move from premisses to a conclusion: to grant the truth of the premisses of a valid deductive argument while denying that, on these conditions, the conclusion would also be true, is to contradict oneself. But the question immediately presents itself. How is one contradicting oneself in such a case? As we saw, the conclusion of a deductive argument need not be an assertorical part of the premisses, so the paradigm case of self-contradiction--of asserting both that A and that not-A--would not seem to be appropriate here if A is the conclusion of a valid deductive argument. Finding out how self-contradiction arises here--and the role that it plays at the foundations of deductive reasoning--will lead us to a closer examination of the notion of deductive validity. What I want to concentrate on, however, is formal validity, that more or less specialised variety of validity which forms one of the cornerstones of the theory of formal logic. My interests in this section centre around formal deductive calculi and I will be

generating the dilemma I have in mind in terms of formal logic. For the sake of simplicity, however, we can restrict the discussion for the most part to propositional (or sentential) logic.

But first, what does it mean for a logic to be formal, to count as a logical calculus? Initially it might be thought that a logic is formal just when it is symbolic, when certain letters, $P_1, P_2, \dots, P_n, \dots$ are used to stand for (say) English propositions and when certain symbols, ' \vee ', '&', ' \supset ', and ' $-$ ', are used to stand for the English words 'or', 'and', 'if...then___' and 'not'. But this conspicuous feature of logical calculi is theoretically unimportant. Symbolism merely provides a sort of shorthand, a compact and perspicuous way of writing down the expressions of the language. What makes a logical calculus formal is, rather, that its primitive symbols--the sentential variables and logical connectives--count as abstractions of English propositions and certain key connectives. The abstraction leaves a residue, and that residue is nothing more nor less than truth-functionality, and it is this which the primitive symbols are used to represent. The complete syntax of the calculus is then built up by means of recursive definitions of well-formed formulas (wff) and proof, this latter notion usually being defined in terms of a set of wffs designated as axioms and a set of primitive rules of inference. In addition, and by way of providing the semantics of the formalised language, an interpretation is set up which consists of a list of functions from the set of sentential variables onto the semantic range (true, false) and of functions defined for each sentential connective from the semantic range onto the semantic range. These functions are designed to exhaust the abstracted 'meaning', the truth-functionality, of the calculus. An interpretation is thus a set of functions; and an

interpretation is acceptable only on the condition that it satisfies two centrally important requirements: (i) no wff will be assigned by these functions both the value 'true' and the value 'false' (non-contradiction); and (ii) every wff will be assigned by these functions either the value 'true' or the value 'false' (bivalence). The resulting calculus is thus formal in the sense that every wff exemplifies a particular truth-functional structure, its logical form. The logical form of any wff, given an interpretation, uniquely determines the truth-value which is assigned to it.

Formal validity is thus a function of logical form. In particular, formal validity is modelled on the truth-functional abstraction of the 'if...then___' construction. (Indeed, all of deductive logic is based on this abstraction--the relation of logical implication. The model of deductive inference is logical implication: one validly draws a deductive inference when, and only when, one shows that the corresponding logical implication is logically true, when the appropriate function assigns 'true' to the implication under every interpretation.) An argument is thus formally valid just in case there is no interpretation of the sentential variables involved which is such that the conjunction of the premisses (the antecedent) is assigned the value 'false'. This much is just a restatement in the language of formal theory of the traditional requirement that one must not be able to legitimately infer a false conclusion from true premisses. But where does the self-contradiction come into the picture?

Consider the following valid argument (an argument expressed in the formal language we are presently assuming):

(5)	$P_1 \supset P_2$
(Therefore)	P_1 P_2 .

Now suppose there were an interpretation, \underline{I} , which would make it possible to say that (5) is formally invalid. Such an interpretation would evaluate the corresponding conditional as 'false', i.e.,

$$(6) \quad \underline{I} [((P_1 \supset P_2) \ \& \ P_1) \supset P_2] = \text{false}.$$

Now (6) would be true just in case \underline{I} assigns 'false' to the consequent of the conditional and 'true' to the antecedent--just in case $\underline{I} [P_2] = \text{false}$ and $\underline{I} [((P_1 \supset P_2) \ \& \ P_1)] = \text{true}$. But it becomes immediately obvious that for there to be such an \underline{I} it would have to assign to either P_1 or P_2 both the values 'true' and 'false'. Thus, to claim both that (5) is a formally valid argument and that the premisses could be true while the conclusion false is to claim that one and the same proposition can be both true and false. And that would be to claim something false about the calculus; but it would not be to contradict oneself. The self-contradiction involved lies deeper in the theory: it is the contradiction which arises from claiming both that \underline{I} is an interpretation (and hence satisfies the requirement of non-contradiction. Thus, the basis of formal validity is derived from a feature of the meta-theory--the definition of 'interpretation'--and is not a notion which is built into the language itself.

All of this is familiar enough, but I bring it out here in order to expose a feature of the semantics of formal deduction: the self-contradiction which might be thought to account for the notion of formal deductive validity--and so to provide a good reason for allowing deductive inferences--is in effect a violation of a rule which is built into the theory of standard, two-valued calculi. Yet, such a contradiction is also an instance of that built-in rule, and indeed it is an instance of the sort to which we would have to appeal to justify the acceptance of that rule. This aspect of formal logic, I want to claim, is the basis for a

dilemma for deduction, a dilemma with far-reaching consequences. But before I develop this dilemma, it is important to consider another aspect of the standard picture of formal deductive validity which has been thought by some to directly conflict with the requirement of epistemic advance.

I am thinking here of the view usually expressed in the opening pages of logic textbooks that truth and formal validity are two notions which must be sharply distinguished. Arguments, it is claimed, are the sorts of things which are formally valid or invalid, and propositions are the sorts of things which are true or false. Furthermore, a formally invalid one may have true premisses. Formal validity is solely a function of logical form: in formal logic we are not concerned with the question of whether the propositions being symbolised are in fact true, or in fact false; all that we need to know is that they must be one or the other.² But what are we to say about arguments which are formally valid but have false premisses? Some have complained that we would never in practice consider an argument with false premisses to be a good argument, even if it is formally valid (vide Hamblin, 232 ff.). Inasmuch as the standard conception of formal validity is based on the formal features of the 'if...then___' relation--a relation which always holds if the antecedent is false--this complaint is an implicit rejection of the basic model for the semantics of formal logic. Given this conception of formal validity, we have examples like the following to contend with:

- (7) If the number of gazelle in Zambia is decreasing,
 then the rate of inflation in England is increasing.

The number of gazelle in Zambia is decreasing.

Therefore: The rate of inflation in England is increasing.

Moreover, there are formally valid arguments with false premisses and false conclusions--again by analogy with the formal features of the 'if...then___' relation. Does the fact that these arguments are formally valid conflict with the requirement of utility?

We might put this question in another way. The utility requirement requires that there be some sort of epistemic advance which results from a piece of reasoning. Now it is fairly common practice among philosophers to treat as the paradigm case of epistemic advance an increase in our stock of factual information. And this was certainly Mill's idea: in the case of the proper form of reasoning, induction, "we proceed from truths which we know, to truths which we did not know; from facts certified by observation, to facts which we have not observed..." (II, 1, 3). There is indeed enough textual evidence to support the view that Mill would not allow as legitimate a piece of reasoning based on premisses not known to be true, or indeed known to be false. But this is surely a cramped view of the nature of reasoning. It would be better to say that the epistemic advance required is new understanding, a fuller grasp of what was previously only partially understood, if understood at all. Sometimes this new understanding can be achieved by an addition to our stock of factual knowledge; but not every advance in our understanding is an addition to our factual knowledge. It might be that we have all the facts before us concerning what Jones did but are not able to say whether what he did was careless, cruel, insulting or evil. And here new understanding consists of a new way of appreciating what is before our eyes. But, then, our question becomes, Where is the epistemic advance, the new understanding, in the case of (7)?

There is, it seems to me, no way of deciding whether there is epistemic advance here until we fill out the context into which (7) is to be placed. (7) appears to be a bad argument because we more or less automatically assume a context which is hostile to it, namely one given by our common sense beliefs about the causes of inflation and the complete lack of impact which the number of gazelle in Zambia has on inflation in England. In this context we have no reason to ponder the value of the reasoning since the implausibility of the premisses vitiates the argument from the start. But in a different context--say that of a story about the odd goings-on in Zambia and the bizarre connection there turns out to be between the number of gazelle and the English economy--in a different context we may have enough of a background provided to accept the premisses and to see the value of the reasoning. In a hospitable context we are in effect being told 'Grant this much...now isn't it reasonable to suppose...?' And if we do grant what the story offers by way of background, it may very well be that (7) gives us a new understanding. In the end we must rely on the context, and our judgement about the value of a piece of reasoning will change from case to case, and from context to context. (We must keep in mind how very reasonable Holmes really was.)

But whatever we decide to say about (7)--or about any other formally valid deductive argument with false premisses (or false premisses and a false conclusion)--there are two points of considerable importance which have been brought out in the preceding discussion: (i) Deciding the question of whether a deductive argument is formally valid is not also deciding the question of whether it is a good argument. A fuller account of what makes an argument a good one will occupy us later, but for the moment we can rely on the two requirements of reasoning we have been

assuming so far. Formal validity is directed toward only one of these requirements; and by satisfying the legitimacy requirement an argument does not automatically satisfy the utility requirement. On the other hand, it is not correct to say, as Quine in one place does, that "'validity' is not to be thought of as a term of praise" (29). It is tempting to say this because the determination of formal validity is essentially a mechanical process. Yet, as we say, the notion of self-contradiction is very closely aligned to the notion of formal validity: invalid arguments are to be avoided because self-contradiction is not exactly to praise him, still, pointing out that he has contradicted himself is surely to censure him. Nonetheless, all that I am insisting on here is that formal deductive validity is not the only evaluative basis for judging the worth of an argument, although, if an argument is deductive, then formal validity is one evaluative basis for such a judgement. (ii) It is surely not the case that we require, in practice, that the premisses of our arguments be true. The value of (7) as a piece of reasoning is not automatically undermined by the fact that the premisses are false (even, on the face of it, absurd). Formal logicians want to hold that arguments with false premisses can be legitimate; but they have a particular model of validity in mind. There is, however, a more important point to be made here. Sometimes arguments with false premisses are good arguments. Those philosophers who object to this suggestion are caught in the grip of an overly-confining model of epistemic advance, one which implies that for an argument to be worthwhile at all it must proceed from truths antecedently accepted to truths previously unknown or in doubt. But there is a vast collection of argument-types which do not follow this pattern: reductio ad absurdum arguments and arguments from counterfactual premisses are obvious examples.

But equally important are arguments from hypothetical cases of the application of concepts, arguments about which we shall have a great deal to say later.

To be sure, a piece of reasoning requires--and as reasoners we presume--a starting place of initial agreement. But agreement as to the truth of a particular set of propositions is only one sort of agreement which is relevant here. One can equally well reason on the basis of an agreement about the initial coherence of a set of propositions which describe a possible state of affairs. We would find nothing at all odd about an argument which rested on the supposition, the imaginary state of affairs, that Wellington lost the battle of Waterloo in 1815. Here we would be interested in a coherent picture of possibility, a self-consistent description of a possible world. The value of the argument would be determined against that background: once the picture was given certain things would follow and certain others would not. But real understanding could result, and the empirical falsity of the premisses would be very much besides the point.

We cannot, then, accuse the formal logician of misrepresenting reasoning simply because on his model of validity arguments with false premisses are sometimes formally valid. Admittedly, that an argument is formally valid is not a sufficient condition for its being a good argument. Formal logic offers a means for satisfying the legitimacy requirement which allows some arguments to be legitimate which might fail to be good arguments in the light of the second requirement. But there is no problem here, we must simply rely on other standards in order to weed out the good from the bad formally valid arguments. There is a problem, however, with the theoretical backing of the notion of formal

validity; and it is to this problem which I now want to turn.

To develop the dilemma I have in mind we need to retrace our steps a bit. Suppose we return to the syntactical notion of proof. To prove any wff B in a logical calculus is to find a finite list of wffs $\underline{A_1}, \underline{A_2}, \dots, \underline{A_n}, \underline{B}$ the last of which we want to prove and each of the rest of which is either an axiom (or an instance of an axiom) or follows from previous wffs in accordance with one of the primitive rules of inference. This stipulative definition of 'proof' is one of the central notions of the theory of syntax. Proofs in logical calculi exhibit two essential features: (i) finiteness, and (ii) a common and unique, antecedently given starting place, namely, the axioms and the rules of inference. Since the word has a wider range of meanings in ordinary language, I shall henceforth speak of 'formal proofs'. The full details of how one actually proceeds to formally prove a wff need not trouble us here, and for simplicity we can assume that the only rule of inference there is is modus ponens (in any event it can be shown that this rule suffices for the proof theory of any standard propositional calculus).

Suppose we now challenge the formal logician to show us why this characterisation of a formal proof is adequate.³ This challenge will take the form of the meta-theoretical question, Is everything which is a formal proof a formally valid argument? To this the logician will respond with a demonstration of soundness: he will show that the axioms are logically true and that modus ponens always and only sanctions moves from logical truths to logical truths, and hence that, given the definition of a formal proof, every formally provable wff is logically true, every formal proof a formally valid argument. This move from the syntactic consequence relation to the semantic correlate is characteristic

of most meta-logical demonstrations, as is the converse move. For, if challenged to show why the formal validity of instances of modus ponens coupled with the logical truth of the axioms should be able to capture, by means of the definition of formal proof, all and only formally provable wffs (i.e., theorems), the logician can respond with the other half of the crucial equivalence between the semantic and the syntactic spheres, that is, a demonstration of semantic completeness. There is, however, obvious circularity of reasoning involved in arguing in this way: one can not both justify one's syntactical notions by showing how they capture the semantic ones, and justify the semantic notions by showing how they capture the syntactic ones.

But the logician can avoid this circularity by claiming that the semantic sphere is prior to the syntactic sphere and hence that the justification ends in the semantic sphere. And this is surely correct, for the syntactical formalisation of logical calculi is, as we saw, an abstraction from natural languages of the truth-functional 'meaning' of the logical connectives. Rules of inference are intended to be syntactic formalisations of patterns of valid argument; and axioms are intended to be syntactic formalisations of patterns of logically true propositions. Thus, the logician will--and to avoid circularity, must--opt to remain in the semantic sphere if called upon to justify the logical calculus. But how is this grounding to take place? Once again, two things have to be shown, that the axioms represent logically true propositions and that the rule of inference is 'truth-preserving', i.e., will never allow an inference from true premisses to a false conclusion.

What we need to do now is to look at the argument used to show that the rule of modus ponens has this essential property of truth-preservation. The inference rule of modus ponens differs from a formally

valid argument of that form--for example (5) above--by virtue of its not being expressed in the syntax of the calculus. Hence:

(8) From $\underline{A} \supset \underline{B}$,
 and \underline{A}
 infer: \underline{B} .

(8) is not a formal proof in the language, it is a pattern of argument-types. And this feature of the rule is made plain by employing the meta-variables \underline{A} and \underline{B} --variables which take wffs of the language as values--and the English words 'from', 'and' and 'infer', and by using the symbol ' \supset ' as a convenient name for the symbol in the language which looks exactly like it.

Now, the argument that the rule given by (8) is truth-preserving proceeds by considering some arbitrary instance of the rule's application. The argument then continues along the lines already given above when the formal validity of (5) was discussed. Suppose we take a look at that argument in the following generalised and condensed form:

(9) Let \underline{I} be an interpretation which is such that $\underline{I} [\underline{A} \supset \underline{B}] = \text{true}$ and $\underline{I} [\underline{A}] = \text{true}$. Suppose, at variance with the rule, that $\underline{I} [\underline{B}] = \text{false}$. If $\underline{I} [\underline{A}] = \text{true}$ and $\underline{I} [\underline{B}] = \text{false}$, then, by the definition of the symbol ' \supset ', it follows that $\underline{I} [\underline{A} \supset \underline{B}] = \text{false}$, contrary to our original assumption. So, our supposition that $\underline{I} [\underline{B}] = \text{false}$ must be wrong, hence (by the defined properties of the interpretation \underline{I}), $\underline{I} [\underline{B}] = \text{true}$. Q.E.D.

Now (9), or something very much like it, is an argument which is familiar to anyone who has gone through some standard meta-theoretic results in detail. But (9) is also a fairly complex deductive argument which consists of several sub-arguments which are not completely stated.

If we were to take the time to state (9) in full, we would discover (i) some suppressed premisses stating the properties of interpretations and the truth-functional characteristics of the symbol ' \supset '; and (ii) sub-arguments which are instances of modus tollens, disjunctive syllogisms of various sorts, De Morgan's laws and modus ponens. But, the point of (9) is to assure us that the form of argument modus ponens is legitimate. Yet, in doing so we have relied (and the examples could be multiplied to make it clear that we must rely) on the legitimacy of particular instances of that same rule. In short: to argue for the legitimacy of a pattern of argument we must employ instances of that pattern. But that pattern is intended to be primitive in the system, that is, it is intended to justify the legitimacy of instances of it.

The difficulty here occurs whenever the formal logician argues about logical calculi. The language in which these arguments are carried out is the meta-language and is invariably a natural language which is augmented with certain symbols (meta-variables and other symbols which occur as autonymous names of symbols in the language of the calculus, or the object-language). The fact--usually credited to R. Carnap--that this meta-language can itself be formalised in the object-language tends to disguise the difficulty which we have brought out. But, although disguised, it remains. For suppose, as could easily be done, the argument given by (9) was formalised in a suitable formal language. Could the various instances of modus ponens which occur in (9), now formalised, be justified by citing the rule of inference modus ponens? Obviously not without circularity. Logicians are not ignorant of this problem and as a result meta-language proofs are never themselves proved as if they were formal proofs in some object-language. The logician relies on the

inherent legitimacy of the particular instances of modus ponens which occur in such proofs--he has no choice. As a consequence, every formal calculus rests on a foundation of informal deductive reasoning, of particular instances of deductive moves which are themselves formalised in the calculus.⁴

Now I want to bring this point alongside the major assumption of formalised logic, the assumption, that is, that the legitimacy of any particular argument or inference is a function of the form that that argument or inference takes. It is this formalist assumption which is embodied in the definition of formal proof: one wff can be inferred from a set of wffs only by virtue of repeated applications of the rules of inference. These rules are intended to be part of the logical superstructure of the calculus, they are the source of the recursive pattern of theoremhood, the foundation upon which the rest of the calculus rests and from which the legitimacy of particular inferences within the calculus springs. Yet, the legitimacy of these rules itself depends on our prior recognition and implicit acceptance of the legitimacy of particular instances of them. These instances may be formalised inferences within the calculus or, as in the argument given by (9) quasi-formal or informal inferences which are outside of the calculus.

The formal logician is thus faced with a dilemma: if he holds to the formalist assumption (provability, in the syntactic sphere, and validity, in the semantic sphere, are functions of form alone), then if he is called upon to justify the primitive rules of inference, he must either embrace the infinite regress which is in the offing, or deny that the 'inferences' his calculus sanctions and generates are legitimate pieces of reasoning. If, on the other hand, the logician abandons the

formalist assumption, then his method of satisfying the legitimacy requirement loses its foundation and the raison d'être of formalising logic appears to be nullified. This, I suggest, is a genuine dilemma for deductive logic because both horns of it are unacceptable. I want now to try to resolve the dilemma by taking a closer look at the sort of reasoning which is exemplified by logical calculi.

Reasoning from Principles

In light of the dilemma for formal deductive logic just advanced we have reason to be suspicious of the project of devising self-sufficient deductive calculi and defining formally 'pure' notions of validity and proof. The obstacle in the way is, ironically, the formalist assumption which motivates the project, the assumption that validity and provability within a deductive calculus must be functions of, and can only be understood with reference to, the logical form of propositions and the truth-functional character of the primitive sentential operators. If the formal logician succeeds in convincing us that his notions of validity and provability are very different from our ordinary notions of valid argument and proof, he also succeeds in making his 'logic' more like mere symbol manipulation. And, conversely, to the degree that the formal logician recognises and admits the conceptual connections between his formal characterisations and informal examples, to that degree will his requirement of formality be abandoned. Both horns of this dilemma require us to give up what has been assumed to be a criterion of adequacy of calculus construction: we seem to be forced either to give up the view that the finished product should reflect how we do reason deductively in informal instances, or the view that the finished product should, because of its rigorous definitions and internal consistency, be able to guarantee its

own legitimacy and fruitfulness as an organon. Now although this dilemma has been developed in terms of the aspirations of the formal logician, the dilemma and these aspirations have analogues in other areas. There is a more general and important mistake about the nature of reasoning which is at the heart of this dilemma, and it is this general mistake which I want to explore in this section.

It is important too to give the formal logician the accolades due to him. For there is little doubt that he can impress us with the systematic power of the calculi he creates. He is like the mathematician in this respect: both can produce genuine surprise and new understanding; they can both move from this proposition to that--and cover a great deal of epistemic distance in the process--without having to plod, step by step, over the ground between the two propositions. When as children we discovered that $2+3$ and $3+2$ gave the same result, and then that $3+4$ and $4+3$ gave the same result, we might have wondered whether there was a pattern here which other pairs of additions would verify. At some point someone might have stepped in and informed us that as a matter of mathematical fact $\underline{n+m}$ and $\underline{m+n}$ give the same result for all natural numbers \underline{n} and \underline{m} . Now there is an inclination to believe that by informing us of this fact we are being told something which is much more fundamental, more nearly like 'knowing how to add', than what we knew about the particular sums $2+3$, $3+2$ and $3+4$, $4+3$. And later the mathematician may wish to reinforce this belief. He may go on to tell us that this is the property of commutativity, that not only is addition of natural numbers and integers commutative, but so is addition of rational numbers; he may tell us that multiplication too has the property of commutativity. He may then go on to tell us that neither subtraction nor division is

commutative; and he may even tell us that addition of transfinite ordinals is not commutative. The mathematician can also give us formulas for solving mathematical problems for which we merely fill in the particular values of the m's and n's. And all of this, and much more, gives him and us a great deal of power over the world of numbers. The formal logician wishes to follow the mathematician's lead--he wants to show us what is fundamental to 'knowing how to reason'. And to do this, he provides us with the appropriate formulas, assuming as he does so that the central logical notions of validity and provability apply to the formulas directly and to the instances of these formulas only indirectly.

So what has gone wrong? One of the points I wish to make in this chapter is that, in a sense, nothing has gone wrong: the formal logician's formulas, like the mathematician's, can be used, and are used, to increase understanding. Yet, our dilemma suggests that these formulas, these generalised principles of reasoning, rely on a recognition that particular instances of them are valid, or are proofs. The mistake is in supposing that bringing forward formulas solves problems, answers questions, resolves confusion, or proves, whereas bringing forward instances can not do these things, or can not do them as well. But here we need to be reminded of Mill's insight: when deductive reasoning is from the universal to the particular, from principles to instances of principles, the legitimacy of the reasoning depends on a prior recognition of the inclusion of the particular in the general. Ultimately, the core of the reasoning is from particular to particular, from case to case; and it is this reasoning which gives the general proposition, the principle, formula or rule its deductive power. It may help me to see the validity of a

particular argument if I am told that it is of the form of modus ponens (it follows that rule). But if I still can not see the validity of the argument I am troubled by, repeating the rule will not necessarily help me out. Here looking at examples of modus ponens arguments, perhaps very simple ones, will be what is needed. But if citing a rule does help me see the validity of the argument in front of me, then that understanding is neither better nor of a different kind than the understanding acquired when I see the validity of my argument by comparing it with other arguments whose validity I can see.

This last point derives from John Wisdom, and is captured in his examples of what he calls 'mother's method'. In Proof and Explanation, Wisdom considers the case of a child who is stumped by a mathematical word-problem: If there are six airlines from England to France, and for each of these six ways of going to France by an airplane and coming back on that same airline, then how many ways are there of going to France on an airplane of a given line and coming back on a different airplane of the same airline? Father attempts to come to the child's aid by giving a covering formula: Whenever there are n things of sort X and for each of these there are n things of sort Y, then there are n times n things of sort Y altogether. This Wisdom calls 'father's method': appealing to a formula, stating a principle or a generalisation, or citing a rule of inference. But mother's method is different, she proceeds from simpler but analogous word-problems: if there are two rooms and in each two books, how many books in all are there?, if there are three different entrances to a building and for each of these three different ways of getting to the main staircase, how many ways are there of entering the building and getting to the main staircase? and so on.

By building up slowly from the simple to the more complicated, mother gets the child to see what the correct solution to his problem is. But is it the case that father's method provides more of a proof than mother's method? It does not, Wisdom answers, for how do we know that father's formula is correct, is what it claims to be: a way of solving all word-problems of a certain sort? We either must appeal to a more inclusive formula to justify father's formula--in which case we just repeat the question about it--or we must resort to mother's method and show why the formula works, by showing how it works in this case and in others.

It might be tempting to think that mother's method is merely heuristically better than father's; heuristically better, but not nearly as rational. In the next chapter I want to bring up considerations which show that this temptation betrays a misconception of mother's method and its power. But for now, to curb this temptation it will suffice to recall a point made in the last section. There we saw that in the end, when the formal logician is required to justify his conceptions of validity and provability, he is forced to employ the sort of reasoning which is the key to mother's method. But reasoning from particular instances is required there not in order to show a child why the formal logician has got it right--although it could be used for that--but to reason about the foundations of the meta-theory of logical calculi. Mother's method is more than a means for bringing out what may be hidden in the formulas, principles or rules; it is a method which provides the power of these generalisations.

But in light of this, are we forced to opt for the second horn of our dilemma for the formal logician--are we forced to question whether

formal logic has a raison d'être at all? What we need to do, I think, is to extract from father's method what is unproblematic and valuable and isolate in it what is the cause of the dilemma. What I want to suggest is that what is unproblematic and valuable in father's method provides the raison d'être of formal logic as well as a wide range of similarly motivated formal systems. To begin, suppose we consider the question, Why is reasoning from formulas, rules, standards, universal propositions, and principles a valuable way of reasoning?

We can begin to answer this question by recalling a Millian point: formulas, rules, standards and principles are summaries of an indefinite number of particular cases, they are ways of organising our thoughts and directing our attention towards features these cases have in common. As a consequence, by citing a formula, rule or principle we are bringing to an argument or an explanation a manifold of particular cases in a manner which picks out the relevant, perhaps structural, features of the particular case we are interested in. In a moment I want to suggest how what is valuable in father's method can be realised; but first it should be noted that those who favour father's method tend to assume that formulas, rules and principles are always pure distillations of all the particular cases they summarise. This assumption gives credence to the view that providing these principles will always expose what is essential about the particular case in question, and expose this in a manner which could not be done by a consideration of analogous cases.

We have already encountered a manifestation of this last assumption about the ultimate priority of father's method: it is, in fact, one of the aspirations of the formal logician. The formalist assumption is the

belief that knowing what validity and provability are really like is knowing how these notions are tied to the notion of the general form of propositions. But this same assumption can be isolated in other claims. It is reflected in mathematical and logical examples where we are inclined to think that, e.g., 'really knowing how to add' requires that one focuses one's attention on principles like 'For any two natural numbers, \underline{m} and \underline{n} , $\underline{m+n} = \underline{n+m}$ '. This assumption is also reflected in the following phenomenon: the inclination to claim that 'really knowing what an X is' is knowing what the definition of 'X' is, being able to give necessary and sufficient conditions for something counting as an X. And lastly, this assumption is reflected in the inclination to claim that 'really' justifying the claim that A was the wrong thing to do on this occasion requires that one show the truth of 'All acts of A-ing are wrong'. What we have here might be called 'Euclid's disease'--an obsession with axioms, formulas, rules and principles. It is an illness which Wittgenstein in the Blue Book has called a "craving for generality", a craving which is accompanied by a "contemptuous attitude towards the particular case" ((2), 18). In the case of accounts of reasoning, this craving for generality is manifested as an obsession with reasoning from principles.

But, it is important to separate the valuable from the obsessional in father's method. Father's method is a method of reasoning and I have claimed that we have two requirements in terms of which a putative method of reasoning can be evaluated. Enough has been said, I think, on the question of legitimacy, so we might turn to the criterion of epistemic advance or utility. What we need to consider is when, and under which conditions, reasoning from principles is valuable because of its capacity

to produce understanding. If we can become clear, from a consideration of the cases of valuable reasoning from principles, how increased understanding may be achieved, then we shall have a remedy for the obsession which the method engenders. For if we have an obsession here then it should be treated as one: the man who has an overwhelming and obsessive fear of being harmed by everything around him, the harmful and the harmless alike, is going to be cured by being shown that there are proper and improper objects of fear, not by being shown that there is nothing to fear. For there are things for which the proper reaction is fear; and what our friend requires is a better grasp of both the harmful and the harmless: he must be cured of the underlying belief that there is no real distinction to be made. Another look at formal logic is in order.

The first and most important thing to notice is that the reliance on principles is built into the whole endeavour of constructing formal calculi. A rule of inference plays the role of a principle in a logical calculus, asserting that every argument with such and such a form is formally valid. This is to be expected since rules of inference are just that, rules: they are intended to have universal application within the calculus. But the axioms of a calculus are also implicit principles, asserting that every proposition which has such and such a logical form is logically true. And indeed, every wff in a logical calculus is in effect a general claim, a formal expression which stands for an infinite number of propositions of that form. (Every well-formed expression in a logical calculus is in fact a formal representation of an infinite number of propositions expressed in natural languages which are symbolised by that expression.) A logical calculus, or formal deduction, then, is

a system of principles, it embodies reasoning from principles par excellence. And at least as far as the formal manipulations within the system are concerned, formal deduction is reasoning both from and to principles.

As for the possibility of genuine epistemic advance by means of the reasoning which takes place within a calculus, we have already noted that this is commonplace. There is a frontier of formal logic; there are advances in the various theories and sub-theories of formal logic; there are problems and there are solutions to these problems. Moreover, we have noted that although the distance which separates axioms from theorems can be bridged by individual deductive moves of great simplicity, in practice this is not, nor need not, be done. Nor is there anything suspicious in the proofs which are given by the formal logician; the jumps he makes from proposition to proposition we can make as soon as we learn his language and his rules. His reasoning is neither unjustifiable nor unproductive; and it is reasoning from principles.

There are two obvious features of the system within which the formal logician's inferences and proofs take place which both explain and justify our faith in his inferences and proofs. These features are (i) the complete systematisation of the elements of the system; the demonstrable existence, that is, of unambiguous connections between propositions coupled with decidable procedures for determining well-formedness and theoremhood; and (ii) a canonical notation which provides explicit extensional reference for all the symbols of the language.

What impact do these two features have on the nature of the reasoning carried on within logical calculi? I want to say that these two features make for an optimal degree of 'determinateness'.⁵ In the

paradigm case of a determinate system--namely a consistent and complete logical calculus--determinateness is a function of the following: complete univocality of predicates (making possible the recursive enumeration of every instance of every 'concept' expressible in the language); and complete systematisation (making it a decidable question, in some systems, whether for any two wffs there is a deductive connection between them). Systems with these properties will be called 'optimally determinate'. This means that every theorem, and indeed every wff, is exactly locatable in the logical space of the system. But then, the picture of a principle as a summary of particular instances now becomes directly applicable. An axiom, for example, captures the form of an infinite number of particular instances of it--there are no borderline or intermediate instances of an axiom. And, in the case of functional logic, there is no ambiguity in the application of a particular predicate--the semantics establishes that some subset of the domain is assigned to each predicate. It is therefore possible to argue within a logical calculus from a consideration of form alone. There are within such a calculus no borderline cases, nor any partial or non-necessary connections between wffs.

Precisely the same extent of systematic determinateness is in evidence in the case of the various theories of mathematics, e.g., the theory of algebraic structures, number theory, probability theory and so on. Here too we might think of those systematised portions of natural science which are physical interpretations of mathematical theories, e.g., mechanics, optical physics, quantum theory and genetics. In all of these cases, reasoning from principles is a highly desirable way of arguing. These theories generally consist of a network of principles, of mathematical or logical axioms, physical laws and

theoretical statements employing highly general, or idealised notions. Reasoning within these systems is often no more than securing a place for one principle or general statement in a network of principles by bringing to light deductive and statistical relationships. In these systems the desirability of arguing from and to principles is a consequence of the high degree of systematisation and the inter-definability of general, theoretical terms. Moreover, in these systems it is often unnecessary, and at times quite inefficient to proceed from a consideration of particular instances (say from particular realisations of physical laws). In the special case of scientific theories, the relation between instance and general statement is obviously much more complex; but within highly systematised areas of natural science, reasoning within the system, and solely in terms of the constituent elements of the system, can be carried out without reference to or consideration of particular instances.

But we really don't have to go as far as quantum mechanics to find examples of highly determinate systems, or at least sub-systems. Although the general field of ordinary language is non-determinate--containing general terms which do not have unambiguous reference and connections between notions which are non-necessary--still there are pockets of ordinary language which are determinate. One obvious case is that of elementary mathematics which forms a determinate system within ordinary language. Another example is the set of kinship terms, like 'mother', 'father', 'sister', 'brother', 'sibling', and the rest. Spatial and temporal relation concepts, geometric concepts and others all group together into compact determinate sub-systems which occur in the open-textured fabric of ordinary language.

When a particular field is highly determinate, the role of the particular instance in reasoning within the system is restricted in practice. It is possible, in logical calculi, mathematical and mathematico-physical theories especially, to rely on generalisations of instances. In logical calculi and certain mathematical theories an emphasis upon form has the consequence that in-system proofs can be and usually are extremely perspicuous. And the striving for what is usually called 'elegance' in logical and mathematical proofs is motivated by a genuine concern: the simpler the proof, the clearer the connection between distinct elements of the system; and the clearer these connections become the greater the chances become of increasing our understanding of the relevant region of the theory or science, or indeed (as in the case of incompleteness or decidability results in logical calculi) of the system as a whole.

In highly systematisable areas of natural science there is an analogous emphasis upon idealisations of concrete examples of certain phenomena, e.g., ideal gases, frictionless surfaces and space-time points and regions. That this sort of systematisation is valuable for a fuller understanding of certain spheres of human knowledge is undeniable. At times greater insight into an area can be achieved by ignoring the infinite variety and concentrating on certain exactly statable features of the manifold of particulars. The direction toward greater determinateness is toward the exactly statable, the regular and the repeatable. Growth in science is measured not merely by the accretion of facts, but by the increased precision in standards and rules. And this increase in precision may often result from moving the standard from one place to another in the system. It is a common phenomenon in the history of

the natural sciences that what was once a standard for ascertaining regularity may in time become an instance of irregularity as measured by a more determinate standard: Galileo discovered the regularity of the swing of a pendulum by using his own pulse; later the irregularity of the heartbeat was ascertained by matching it up with timepieces which employed pendulums.

What I want to claim, then, is that the value of reasoning from principles, from generalisations and idealisations, is readily apparent when that reasoning takes place in a particular sort of conceptual environment, namely within a system of sufficient determinateness. Within such systems, reasoning is successful if it provides a fuller understanding of the relevant principles, generalisations and idealisations. But, the field must be susceptible to the sort of systematisation which provides the suitable environment for reasoning from, and to, principles. To see the value of father's method, then, is to see where it can be employed to produce genuine understanding. And as a consequence, to see what sort of environment is required for the application of father's method is to see as well where that method is not appropriate, yields confusion and prevents understanding. Consider Wisdom's example of the child and his word-problem once again.

In Wisdom's example the child is puzzled over a very specific mathematical problem. What he needs is an understanding of the use of multiplication in this sort of case. But he also needs to see where the numerical data is hidden, what relationships are important, and what is relevant to the question and what is not. Finding one's way around the problem of the airplanes and the airlines requires the same eye for clues, the same ability to discern pattern and recognise

numerical concepts required for finding one's way around any mathematical problem. But the patterns the child can not see, we may have less trouble in seeing. And the patterns which we may not be able to see, the mathematician may find obvious. It is to be expected that what it takes to see a mathematical problem in a way which brings the data and the relevant relationships to light will depend on how much of the surrounding patterns and concepts one is familiar with. What father offers as a clue to the solution to the child's problem is a clue which requires preliminary experience at finding one's way around in mathematics. Father's solution fits into a system of mathematical principles and concepts in terms of which it is a solution. Mother's method, on the other hand, reveals the same pattern but in a manner which does not rely on as much of an understanding of the system as is presupposed--indeed contained in--the formula which father provides. When as children we found our way to the candy store, we probably began by going 'that way' at the corner where the butcher shop was and 'this way' where the post office was. But when we acquired an understanding of the concepts of left and right we could use these concepts to inform our friends where the candy store was. Still later, when our immediate environment was seen as a pattern of streets related in such and such a way to other streets in our city, and then related in such and such a way to other cities which were in turn related to each other in countries and continents, we acquired an understanding of 'North' and 'South' and the rest.

But, if father's method incorporates reasoning which depends on an interrelated system of principles and concepts, mother's method incorporates reasoning which provides an understanding without which that system of principles and concepts would be incomprehensible. As we saw in the last

section, the principles and concepts which are foundational to logical meta-theory rest on an understanding of instances, and ultimately informal instances of valid arguments and proofs. So, part of the therapy for the obsession for reasoning from principles is just the reminder that the value of this sort of reasoning depends on an understanding of the principles and concepts employed and ultimately this understanding can only result from an understanding of particular cases.

Another manifestation of the "craving for generality" is important here. It is the obsession with definitions, usually expressed as the claim that a real understanding of a general notion can only come from a definition of that notion. Wisdom gives us this example of the application of mother's method to the problem of understanding a general notion:

A child asks, 'What is a greyhound?' His father replies, 'A greyhound is a dog of a certain sort.' 'I know', says the child, 'but what sort?' 'Well', his father says, 'a greyhound is a dog in which the power to weight ratio...' But his mother interrupts. 'Look', she says, 'that's a greyhound, and you remember your uncle's dog, Entry Badge, well that was a greyhound. But now that', she says, pointing to a Borzoi, 'is not a greyhound, and even that,' she says, pointing to a whippet, 'is not'. ((2), 69)

Here father's approach is to attempt a definition, and he tries two of them. The first is too vague--it leaves open exactly what the child wants settled--and the second introduces considerations which would be informative only to someone who already had some idea what a greyhound looked like. Mother's method, as before, employs examples, and importantly, examples of greyhounds as well as examples of near-greyhounds. A greyhound is a certain sort of dog; but it is bigger than a Pomeranian and smaller than a Newfoundland. If you know what a Borzoi looks like, then think of that animal slightly smaller and with

less fur. It's bigger than a whippet and doesn't have the long hair around the ears and tail like a Saluki...and so on. It is not that what father says is false; the first definition he gives is, if anything, too true; and the second, like the formula he gave in the other example, is not wrong, it merely requires more experience with the subject than the child has. Nor is what mother says a substitute for a definition, something which is tailor-made for this particular inquisitive child. Rather, what mother provides is a way of understanding what a greyhound is, as dogs go, which relies on similarities and differences being brought to the child's attention. And this procedure has a more general application than it might at first appear to have; in the next chapter we shall see the procedure at work on questions like 'What is negligence?' and 'What does it mean to be careless?'.

To be sure, some general notions have very clear boundaries, some terms can be defined with precision, and principles for the correct application of such terms can be specified. And these definitions and principles may put all that needs to be known in front of one's nose. But, does real understanding of a general notion require a definition? When the mathematician defines a prime number as 'a number having no integral factor except itself and unity' what he says is exactly right. Giving this definition may be all that is required for an understanding of what a prime number is. But if it does give us all we need to know, it is clear that we already must have an understanding of the constituent notions which are employed in it. Divisibility, for example, is a notion which is used in the definition, and it in turn may be defined in terms of other mathematical notions. But then, understanding the notion of divisibility must then rest on a prior understanding of the constituent

notions employed in its definition. We could keep this up for some time, but this much should be obvious: at some stage our understanding of at least some of the constituent notions employed in these definitions must be an understanding of particular examples of division, of a factor, of a natural number, or of a prime. It is likely that we can see that the definition of a prime number given above is exactly right because we know what it is for a number to be divisible only by itself and 1. And we have an understanding of what it means for a number to have this property when we look at the numbers 1, 2, 3, 5, 7, 11, So, even when definitions are exactly right, our understanding of them, our realisation that they are exactly right, depends on an understanding of constituent notions employed in the definitions which in turn, ultimately relies on an understanding of particular cases.

I have been suggesting so far that what is valuable in reasoning from principles and providing exact definitions is an increased understanding. Father's method relies on the existence of highly determinate systems of principles and general concepts, his method constitutes reasoning within such systems. Now the dilemma which was posed in the last section can be restated for the other examples of determinate systems we have considered. And the lesson to be learned in each case would be the same: the features of these systems which guarantee the power and perspicuity of the reasoning which can be carried out within them can be realised because we have a prior understanding of instances of at least some of the principles, and of particular cases of at least some of the concepts, employed. But I have only considered those areas of knowledge which are amenable to this sort of systematisation. A major symptom of Euclid's disease is the belief that systematising is a

pre-condition for complete or genuine understanding. By looking more closely at this symptom we can begin to see the intrinsic value and the indispensability of reasoning from particulars.

Footnotes

Chapter II: Reasoning from Principles

(Page 32)

¹It is interesting to compare Aristotle and Cicero on reasoning with the account given of proof by the Classical Sceptics. For Sextus Empiricus, for example, a proof must be an argument ('a collection of premisses and a conclusion') which is conclusive (deductively valid) and true (sound), and which is such that the conclusion is non-evident (adēlon) though its plausibility is discovered through the power of the premisses (Vol. I, 310 ff.). Even though Sextus Empiricus was interested in showing that such proofs are impossible (and so there is some reason to suspect ulterior motives for the stringency of the account), still both requirements for reasoning are embodied here. The demand that a proof be both valid and sound is the demand for legitimacy; and the (somewhat mysterious) demand that the conclusion be non-evident is the demand for epistemic advance. Interestingly enough, however, Sextus Empiricus has another requirement which in effect links these two other requirements: the conclusion must constitute new information which is brought to light by (and only by) the 'power of the premisses'. This last requirement allows Sextus Empiricus to set forth a quick argument against the possibility of proof: the power of the premisses is never sufficient to establish a non-evident conclusion, for if this were possible then what is non-evident could be demonstrated to be evident; but that is impossible. This argument is, moreover, closely linked to another which also depends on this third requirement: every demonstrative argument--i.e. every deductive argument which purports to demonstrate its conclusion--is in fact an example of the fallacy of petitio principii. In the next few pages I shall be considering Mill's version of this claim (which was obviously borrowed in toto from Sextus Empiricus). Mill's argument, I shall claim, depends on the misleading suggestion that in a valid deductive argument the conclusion must be 'contained in' the set of premisses. What lead Mill astray, it seems to me, might be thought of as an implicit adoption of Sextus Empiricus's third requirement. So, in what follows, I hope to show why both the 'contained in' metaphor and Sextus Empiricus's third requirement can mislead.

(Page 48)

²This point is often put in terms of the hypotheticality of deductive inference: we do not need to know whether any of the constituent propositions of a deductive argument are true, A. Flew (for example) has claimed, because "we are not actually saying anything about the truth or falsity of those constituent propositions. It is all hypothetical" (11). To this Hamblin has objected: "A real argument has real premisses and conclusion, not hypothetical ones" (233). Yet, in this context the opposite of 'hypothetical' is not 'real' but 'known to be factually true (or false)'. To argue from a hypothesis or a supposition is really to argue, and to argue hypothetically is not necessarily to pretend to argue (or to argue half-heartedly, or whatever it is Hamblin has in mind); it is rather to argue from a basis other than agreement as to the truth of the premisses.

(Page 53)

³The line of argument I am going to offer here can also be put in terms of the challenge to 'justify deduction', a challenge intended to parallel that made with respect to induction. It has been suggested (most recently by S. Haack) that Hume's dilemma for induction has an analogue in the case of deduction: we cannot justify deduction inductively since that would at best show that usually when the premisses of a deductive argument are true, the conclusion is true as well (and that would be too weak a justification); nor can we justify deduction deductively, since such a justification would be circular. Now the demand to justify deduction is obviously the demand to show that deductive reasoning--for example, the inference rule modus ponens--meets the legitimacy requirement of reasoning. But, although I shall not in the text consider separately the problem of justifying deduction, the solution to the second dilemma for deduction which I shall propose below counts as well as a solution to this problem: We cannot hope to justify deduction either deductively or inductively; in the end the justification of deductive inference rules can only arise out of a non-deductive, non-inductive consideration of the instances of those rules. The objection that a reflective justification of deduction would be, like an inductive justification, too weak, will be at least implicitly met in Chapter V when I consider the question of the legitimacy of reflective reasoning.

(Page 57)

⁴Formal purists might object that a logical calculus can be treated as a wholly uninterpreted system for the manipulation of certain symbols, and when it is so treated there is no dependence upon, since there is no connection to, informal deductive arguments. If the formal logician wishes to disassociate his symbol manipulation from the activity of reasoning, then he cannot be touched by the dilemma I am here outlining. Justification for him ends with 'that's the way these symbols are defined'. There is nothing particularly incoherent in this view; it is just extremely uninteresting (and quite out of line with the historical development of formal logic).

(Page 66)

⁵This term is adapted from Abraham Edel's notion of 'field determinateness' (vide (1) and (2) passim). Although my use of the notion of determinateness differs from his, there is nonetheless a close analogy which is noteworthy. For example, Edel claims that the field of ethics would become determinate if the following conditions could be met: (i) if the concept of a moral phenomenon could be firmed up in order to mark off an area of inquiry--say observed behaviour--which was distinctly moral; (ii) if the central moral concepts--good and evil, right and wrong--could be linked to this delimited area of moral phenomena so that it would be possible to provide necessary and sufficient conditions for something or some act being an instance of one or another moral concept; and (iii) if there turned out to be sufficient regularity to allow for empirical generalisation, for systematisation which is suitable, from a logical point of view, for empirical verification.

Now my use of the term is designed primarily--as it will become apparent below--to designate a kind of conceptual environment in terms

of which deductive reasoning manifests its inherent usefulness: the appropriate conceptual environment for deductive reasoning. I shall later also use the notion in the context of inductive reasoning. With Edel, I will speak of determinateness primarily as a systematic property and, specifically, as a property of concepts which fit into determinate systems or sub-systems. Any concept, I should hastily add, can be made determinate simply by setting out strict necessary and sufficient conditions for its application. But some concepts resist this treatment, i.e., any set of conditions provided will be met with counterexamples.

CHAPTER III

REASONING FROM PARTICULARS

Having seen some of the reasons for the attractiveness of the purely deductive model of reasoning from principles in the last chapter, here that model will be subjected to two challenges. The first is aimed primarily at the notion of justification, a notion which is for many philosophers the same as the notion of subsumption under a general principle or rule. The second, and more extensive challenge has to do with the fact that there are non-determinate areas of human knowledge, areas for which reasoning from principles is not what we might call the 'working logic'--the mode of reasoning which solves the problems, answers the questions and resolves the perplexities which are characteristic of the area of knowledge.¹ The non-determinate system which gets the most attention here is the common law. The law will play an important role for us in subsequent discussions, both as a source of examples and as a sort of therapeutic comparison, a reminder of how essential non-deductive and non-inductive reasoning is. It would not be too misleading to say that when most philosophers consider the nature of reasoning, of being rational, somewhere not too far back in their minds is the picture of the uncluttered and austere, the simple yet powerful, deductive logic machine. I would like to replace that idol with the workings of the legal forum.

The 'End' of Justification

The dilemma for formal deductive logic posed in the previous chapter was shown there to rest on a general mistake about the nature of reasoning. The mistake has to do with a misapprehension about the source of the legitimacy of deductive inferences and arguments. And the solution to the dilemma arises out of the realisation that formal deductive calculi cannot internally guarantee their own legitimacy and fruitfulness as systems of reasoning; ultimately, when the question of legitimacy is seriously raised and pursued, the source of this legitimacy is found to lie in informal and particular deductive inferences. But the formalist assumption, in its various manifestations, is an attractive assumption. It is attractive for at least two reasons: The first is that within highly determinate systems, like logical calculi, mathematical theories and certain systematisable areas of natural science, we do in fact proceed in the way the formalist assumption suggests: when called upon to justify a particular inference, or a particular instance of a natural law, we employ father's method, we appeal to the rule of inference, the principle, the natural law. I tried to show in the last section of the previous chapter why justification, as well as epistemic advance, can result from reasoning from principles within highly determinate systems. And if we take such determinate systems to be paradigmatic of every area of knowledge where reasoning is required, then it is tempting to assume that reasoning from principles is the proper mode for ascertaining the legitimacy of claims made in any area. The second reason why the formalist assumption is attractive is that it appears to bring the request for justification to an end by exposing what are thought to be the necessary pre-conditions for the validity of inferences or the truth

of scientific claims. In the next section I want to deal with the first of these reasons for the attractiveness of father's method by considering in some detail the nature of reasoning within a non-determinate system, the common law. Here however I want to consider the second of these reasons and to look at the source of the legitimacy of deductive reasoning more closely.

In the case of formal deductive logic, the formalist assumption asserts that the legitimacy of any particular inference is a function of the form of that inference, and hence that a particular inference is formally valid if, and only if, it is an instance of some rule of inference. The dilemma posed for formal deductive logic grew out of the fact that, when called upon to justify these rules of inference, the formal logician must rely on the validity of particular instances of them. The result is that, if the formalist assumption were true, there could be no, non-circular justification of the legitimacy of deductive inferences. This problem has been expressed in a different way by Lewis Carroll in his "What the Tortoise said to Achilles". In Carroll's tale Achilles and the Tortoise are considering the argument 'If A and B, then Z'. The Tortoise poses the question of why, if he accepts the truth of A and B, he is forced to accept the truth of Z. Achilles dutifully responds with (C) 'If A and B are true, Z must be true'. C here is in effect a rule of inference which, were it accepted, would force the Tortoise to accept Z. The Tortoise accepts C but notes that it must now constitute one of the premisses in the argument, one of the necessary steps on the way to the acceptance of Z. The Tortoise then challenges Achilles to show him why he is now forced by logic to accept Z. Achilles responds with (D) 'If you accept A and B and C, you must accept Z', yet another

rule of inference. But the Tortoise correctly notes that D must now be another necessary step on the way to accepting Z, and asks Achilles why, if he accepts A, B, C and D, he is logically compelled to accept Z. Carroll thus delightfully offers us another race between Achilles and the Tortoise, the race whose finish line is the acceptance of Z. But it is a race which can never finish as long as Achilles sticks to the assumption that the only way one can be logically compelled to accept the conclusion of an argument is by accepting the truth of a covering rule of inference.

The Tortoise's challenge to Achilles is analogous to our challenge to the formal logician in the last chapter: what reason can you give me for thinking that the epistemic advance provided by the conclusion of a valid argument is justifiable? Moreover, the initial move made by both Achilles and the formal logician is the same: to appeal to a rule of inference. In the case of Carroll's tale, the appeal to a rule is maintained in the face of each of the Tortoise's challenges; and the result is infinite regress. I have suggested that the solution to the dilemma which is embodied in the challenge to the formal logician consists of rejecting the formalist assumption. In terms of the on-going debate between the Tortoise and Achilles the solution is analogous: Achilles must realise that the procedure of justification he is relying on is ineffectual against the Tortoise's challenges; his justification is proceeding in the wrong direction. Instead of proceeding upward, so to speak, to more and more complicated rules, each of which entails the inference which preceded it, Achilles should proceed horizontally and bring up analogous cases of the first inference which troubled the Tortoise. Achilles' procedure is in fact analogous to that of the doctor

in Moliere's comedy Le Malade Imaginaire who, when asked to explain why opium produces sleep when administered, replied that it has a vis dormiliva, a soporific quality. The circularity in this 'explanation' parallels that involved in the formal logician's justification of the validity of an instance of modus ponens. And in the case of the justification, the escape from explicit circularity, or infinite regress, lies in an exploration of analogous cases, and a denial of the pre-eminence of the hierarchical pattern of justification by entailment.

Yet, it must always be borne in mind that all reasoning about the justification of reasoning is inherently circular: in so far as one is being rational, one is utilising forms of reasoning whose justifiability is the object of the investigation. But this sort of implicit circularity must be sharply distinguished from the explicit circularity of arguing for the legitimacy of forms of argument by citing rules whose legitimacy relies on those forms of argument. Indeed, the inherent circularity of reasoning about the justification of reasoning loses its viciousness if we proceed by bringing up examples of justified inferences, if, that is, we show or exhibit the clear cases. One might think here of Wittgenstein's somewhat cryptic remark at the end of his Tractatus, the propositions of the Tractatus are indeed nonsensical (nonsensical pseudo-propositions to be precise). But this does not vitiate the project of the work since these propositions are also elucidations, they show something of importance about the nature of language and the world (cf. Bickenbach). Similarly, the activity of bringing up cases of valid inferences--to elucidate what deductive validity is like--is a procedure which escapes explicit circularity while, as we shall see later, remaining wholly rational.

But the citing of examples has traditionally been viewed as having a purely heuristic value, and being ancillary to genuine justification which consists in exposing the rules, principles and standards which are the necessary pre-conditions for the legitimacy of inferences or the truth of scientific claims. We saw in Chapter I how this conception of justification has monopolised the field in the case of the justification of moral assertions, and we saw in the previous chapter the far-reaching ramifications of this view in the areas of logic and science. The appeal of this conception of justification rests on a very simple point of logic: if the justification of a proposition Q can be achieved by discovering another proposition P which entails it, then the end of justification consists in the discovery of an ultimately general proposition A which entails P, as well as all other, intermediate propositions between it and P, and hence, eventually Q as well. But how is this the 'end' of justification? One standard response is that no more justification can be required since no more can be given. But why should we accept this claim?

Herbert Feigl in his "De Principiis Non Disputandum...?" has sought to underwrite the view that justification can reasonably be said to end. He has argued that there are really two forms of justification, validation and vindication. Validation necessarily involves an appeal to principles, rules or standards, and so, logically, reaches a point beyond which no further validation is possible. Still, at this point, the point at which ultimate presuppositions have been reached, we can satisfy the request for more justification, although the justification given is of a radically different sort: it is the pragmatic vindication of these principles, an answer to the question, Why should we adopt just these principles?

According to Feigl, a pragmatic justification is an attempt to show that principles serve as means toward some end whose desirability is presupposed. Justification at the level of ultimate principles is not viciously or explicitly circular, Feigl argues, because the question is no longer one of validation; it is rather a question of showing of these principles that "their acceptance as regulative standards is an indispensable prerequisite for the fulfillment of certain ideal requirements" (Ibid., 144). To be sure, claims about these ideal requirements and means-ends relations are themselves claims which require justification in the sense of validation; but the circle is broken because some of the requests for justification can be met by vindications. Hence, justification only appears to be circular if one equivocates between validation and vindication. Furthermore, justification qua validation and justification qua vindication do end, the former in principles and the latter in purposes, needs or ideal requirements.

As we shall see in the next chapter, the notion of justification by vindication has been employed to tackle the question of the justification of induction (and M. Dummett has suggested that a similar line might be taken in the case of deduction (Dummett, 17 ff.)). Yet the distinction between validation and vindication--or at least the way in which Feigl makes out this distinction--obviously prejudices the question of whether an appeal to principles constitutes the only proper way of non-pragmatically justifying particular claims or inferences. Validation just is appeal to principles, rules and natural laws. But, it is not always the case that when justification in the sense of validation ends with the assertion of a principle the only question which can then be asked is, What purposes or requirements are satisfied by this principle?

Consider the example of modus ponens again. If the formal logician claims that by citing this rule of inference he has gone as far as he can go to satisfy the requirement of validating the legitimacy of some instance of modus ponens, then, if we are concerned about the validity of the rule itself, we would not be at all satisfied with the supposed vindicating reply--Since logic is concerned with what is required in order to be rational, we must accept the validity of modus ponens. Rather, our problem is just that of finding reasons for thinking that accepting the validity of modus ponens is a rational thing to do. The formal logician may attempt to show us that modus ponens can be justified--validated in terms of a more primary rule, namely the principle of non-contradiction. And at this point he seems to be on firmer ground since the question of whether this principle is justifiable, in the sense of being capable of vindication, would appear to be self-answering. But if we are happily convinced that the principle of non-contradiction can be vindicated by setting the principle up against the intuitive requirements of what it means to be rational, we should nonetheless be alive to the fact that our being convinced is just our being aware of the rationality of particular instances of the principle. The principle is secure as a canon of rationality for at least one good reason: we are all quite aware that it is not rational to claim, say, that you will find the butter in the pantry and you won't find the butter in the pantry, that it is both night and day at this spot on the globe, that something is red and green all over, that I have a brother who is not my sibling, and on and on.²

The dilemma I posed for deduction is a dilemma about justification, the justification of the legitimacy of deductive inferences. But it is

not a dilemma which can be resolved by retaining the formalist assumption--and so, Feigl's conception of justification as validation--and injecting at the appropriate place the notion of justification as vindication. The notion of a kind of justification which counts as a pragmatic vindication is a valuable one; but the manner in which Feigl sets out the distinction is misleading because artificially strong. For Feigl maintains that questions about vindication can only be properly asked of ultimate principles--principles of non-contradiction, induction, verifiability and justice being the examples he has in mind. In so far as it does seem true to say that general principles or rational policies are the sorts of things standing in need of pragmatic justification, Feigl is right to reserve vindication for such general assertions. Yet since, as I have been suggesting, the legitimacy, the credibility and the rationality of these principles is just the legitimacy, credibility and rationality of instances of them, validation and vindication in the end would seem to merge. In particular, justifying an instance of a principle will involve a consideration of the rational, epistemological or moral ends to which the instance is seen as a means.

There is no doubt that Feigl's distinction as it stands does seem to be appropriate to questions about the justification of propositions within the boundaries of highly determinate systems as opposed to questions about the justification of those systems themselves. It is, I have argued, reasonable to argue from principles within such systems--determinate systems have the kind of conceptual environment which makes reasoning from principles appropriate--and such reasoning is at the heart of Feigl's notion of validation. It should be noted that questions of vindication in such contexts can only be properly asked

of the system as a whole. Moreover, if the issue is one of pragmatically justifying a determinate system, it is also appropriate to ask whether the internal logical structure of the system satisfies the rational, epistemological or moral ends we have in mind. Thus, the whole procedure of justification as validation would seem to require vindication, and a separate vindication for the different fields in which that justificatory procedure is applied. Such a vindication I have already attempted to sketch out in the case of formal calculi. But in an important sense, the purposes and goals which a particular determinate system is intended to satisfy are suffused throughout the system: questions of pragmatic justification, even if initially asked of principles, are generally answered in terms of the effect which instances of them have throughout the system. Nelson Goodman was thus on the right track when he claimed that, in deductive and inductive logic, justification takes the form of a mutual adjustment between rules and accepted inferences and that this is an on-going process (62-66). It is, it seems to me, merely a remnant of the rationalism which Feigl claims to eschew which is the basis for the insistence that questions about pragmatic justification can only be asked of the ultimate principles which are embodied in the system. But then, it is just another symptom of what I have called Euclid's disease to see the justification of the 'theorems' as requiring one sort of procedure and a justification of the 'axioms' another, entirely different. The remedy here is to realise that the ultimate principles require, and are susceptible to, the same sort of justification as the instances of the principles. Thus, Feigl's attempt to avoid the perplexing question of why it is rational to suppose that justification ends with ultimate principles fails to satisfy our qualms about the

inadequacies of the hierarchical, or Euclidean, model of justification.

But if, as I have been recommending, we reject the view that the hierarchical or Euclidean model is the only possible model of justification, then where are we to locate the 'end' of justification? I have been insisting that, in the end, questions about the legitimacy of deductive inferences are resolved at the level of particular, and informal, deductive inferences. In the end, reasoning about particulars is the heart of justification, as it is the heart of explanation and proof. But then, does justification end with a consideration of particulars? A short answer to this--as a prelude to a more complete answer in Chapter V--is: yes and no. Yes, if we are taking the 'end of justification' to mean that range of instances out of which justification inevitably arises; but no if we are taking that phrase to refer to a point beyond which no more justification can be legitimately requested. The request for justification, I want to suggest, can always be legitimately made, and, if the object of the request is justifiable at all, can always be answered. But not answered in such a way as would preclude the possibility of more justification being forthcoming. In this sense, justification has no end. There is no absolute ground of justification, no set of ultimate principles which themselves only justify but can never be justified. Nor is there a set of ultimate particulars (e.g. sense-data reports) which has these properties. But there is a source of justification, namely particular cases, instances of principles and clear examples.

Reasoning from Particulars and the Law

As we have noted, reasoning from principles in highly determinate systems is made possible both by the internal logical structure of such

systems and by the nature of the concepts employed therein. Suppose we call the principles (natural laws, rules, formulas, standards and other such generalisations) and the concepts which are incorporated into such systems determinate principles and concepts. A concept, C, is determinate in a system if, and only if, there is a determinate principle of that system which asserts that all and only the things (strings of symbols, events or other phenomena of the field) which possess each of a finite set of properties P_1, P_2, \dots, P_n are C's. Moreover, a system may be said to be closed when, for every concept in that system, each of the defining concepts P_1, P_2, \dots, P_n is itself a determinate concept in that system. In practice, of course, some concepts are treated as primitive or initially undefined when one sets about constructing a system of this sort. Thus, for example, the concepts of natural number, zero and successor are primitive concepts in the axiomatic theory of natural numbers which is built upon Peano's axioms. But it is nonetheless possible, although with obvious circularity, to define the primitive concepts in terms of other concepts in systems of this sort. This property of mutual interdefinability provides such systems with their distinctive internal cohesiveness. When the concepts of a system are not exhaustively interdefinable in this manner then the system is open, and so not optimally determinate. We shall say, however, that every instance of every concept in a closed determinate system is uniquely recoverable in the following sense: In the case of the concept of prime number, for example, it is the case that although there is an infinite number of instances of this concept, no instance of a prime number is also an instance of another concept in the system of formal number theory which is incompatible with the concept of a prime. That is, if n is a prime

and also an instance of another numerical concept C (e.g., square root), then for none of the defining properties of a prime is the negation of that property one of the defining properties of C, and vice versa.

The paradigm examples of closed determinate systems are logical calculi and mathematical theories. Yet there has been an on-going debate in the philosophy of science over the question of whether defferent areas of natural science can profitably be turned into closed determinate systems. The question, moreover, changes character if asked first of physics, then of biology and then of psychology and sociology. (But the question is not clear-cut even in the case of physics: as Einstein once claimed, "So far as the laws of mathematics refer to reality, they are not certain; and so far as they are certain, they do not refer to reality" (quoted by Waismann, 17).) In some cases, concepts introduced to account for the fundamental difference between one science and a more determinate one (sciences which have, as it were, common boundaries) have been eliminated in favour of other concepts which provide the interface between the two sciences. Thus, up until the early 1800's it was thought that the concept of a vital force was required to account for the synthesis of organic compounds and to support the view that organic compounds could arise only from living sources. The emergence of biochemistry was made possible by the German chemist Wöhler who in 1828 showed that the inorganic salt ammonium cyanate can be converted to the organic compound urea by heat alone. Nonetheless, certain concepts in biology have resisted definitions which would make them determinate with respect to chemistry or physics. The concept of living material is the most obvious example. There is little doubt that a cat or a jellyfish is living, and that a lump of barbituric acid

or a chair is non-living matter. But a bacteriophage virus is just as clearly a borderline case. Of course, every science has a partial skeleton around which some flesh clings, since, for example, every science incorporates statistical machinery. But a 'soft' science like sociology just does not have the same sort of internal structure that physics has, and some philosophers have expressed doubts whether sociology ever could have such a structure.

But what we need at this point is a field of questions and problems that require reasoning which is intermediate between the fields we have just considered--which themselves constitute a range of fields from logical calculi to the soft sciences--and fields like morality or aesthetics. It seems to me that we have in the common law just the sort of intermediate area we require. The common law is of interest for another reason--it can be viewed as an analogue of morality. There are analogies and disanalogies between legal, or judicial reasoning and moral reasoning, and there are parallels between jurisprudential accounts of judicial reasoning and philosophical accounts of moral reasoning.

It is, I think, both important and intrinsically interesting to consider the example of the law in some detail. It is usually the case that a consideration of the law reveals the results of serious philosophical inquiry which has produced philosophically significant conclusions. These conclusions amount to jurisprudential decisions demanded by the practical constraints of maintaining a viable legal order. While the decisions made on what are philosophical questions about proof, reasoning and evidence are not always philosophically unobjectionable, a consideration of these decisions and accompanying obiter dicta is always worth the effort. What I want to do here is to

introduce two meta-legal issues relevant to our interest in reasoning from principles. These two issues are (i) the nature of a legal system; and (ii) the character of reasoning within the common law. There are, of course, various ways in which these two issues could be discussed--there is the purely jurisprudential approach, the everyday, lawyer-in-court approach, the historical approach, and others. But for this present discussion I want to keep to a fairly high level of abstraction and consider these two issues in the light of the predominant jurisprudential schools. I do this because the mistakes and distortions which are packed into two of the major traditions are not only familiar but illuminating.

When jurists consider the notion of a legal system they tend to see and describe two different phenomena--the propositions of law, and the activity of law-makers. Thus, one tradition, called 'Analytical Jurisprudence' (which in its extreme form, I shall call 'Formalism') tends to see a legal system strictly in terms of what, quite literally, is 'on the books'. The concentration here is on the laws themselves, the statutes of the various Codes and the precedents of the Law Reports. (It should be remarked here that I shall only be concerned with the English tradition in law and not with the Civil Law tradition which lacks a common law background.) These legal propositions have been thought to reflect a fundamental unity of structure, or system. And the unity and system is commonly thought to take on a hierarchical arrangement such that the legal validity, the sanction or 'justiciability' of particular statutes or precedents relies on other statutes and legal principles (e.g. equity) which are higher up the hierarchy. Even within a single area of law--say, the law of torts--the arrangement is thought to be hierarchical with certain

implicit principles--reasonable care, the right to property and the like--being higher up the systematic hierarchy than specific injunctions, or precedent decisions, concerning negligent acts or nuisances. Suggestions for the highest level of this pyramid include the classic Austinian claim that the law grows out of a command issued by a sovereign or other authority, and the Kelsenian position that the law rests upon an 'apex norm' (e.g. 'The constitution of the legal system of...ought to be obeyed'). Alternatively, the unity of a legal system has been argued to be founded upon certain primitive legal concepts such as right, privilege, power and immunity (which are then correlated with duty, no-right, liability and disability). Implicit in the enterprise of Analytical Jurisprudence, and the central thesis of Formalism, is the view that the whole body of law in a legal system is self-consistent and analysable into a fairly small number of basic legal norms or concepts.

Another jurisprudential tradition--one which has grown out of a rejection of Formalism--emphasises the facts of the everyday functioning of courts. It sees a legal system in terms of what legislators, lawyers and judges do more than in what they say, and hence, what is codified into law. This approach (the 'Functional', 'Sociological', or 'Realist' school), views law in the context of the prevailing social, economic and even psychological currents. Embodied in this view is a sceptical attitude towards the 'purity' of legal concepts and an explicit denial of the supposed ineluctability of logical connections between statutes and precedents and instant cases. A legal system is not, on this view, a body of laws; it is a collection of sociologically significant facts--namely, legislative and judicial decisions. Thus, Jerome Frank, a

particularly hard-nosed Realist, has accused the Formalist of legal fundamentalism, of worshipping rules and exaggerating the element of certainty in the law. At one point, Frank even accuses the Formalists of failing to realise that a judge's digestion may be just as potent a factor in influencing his decisions as the rules, or reasons, he cites in his decisions (Frank, 41, 57 and 69). De-emphasising the systematic character of the law (often to the extent of viewing it as the expression of the emotions of legislators and judges), the Realist substitutes a picture of an institution firmly placed in the realities of social, economic, political and other extra-legal pressures. Figuring centrally in this picture is the element of choice on the part of the judge: he is no logician, nor is it at all likely that he is a Solomon; he is rather a maker of law whose decisions reflect (and indeed, can be justified by) social, economic and political conditions and circumstances. In a word, the theme advanced by the Realist is that there is no system to the law, there is no logic to the decisions of judges. The law should be studied as a social and psychological phenomenon, not as a system of propositions.

Now these two jurisprudential schools of thought fairly clearly represent polar extremes. Both views, moreover, are exaggerations of the nature of a system of law. The view that the collection of legal propositions of any legal system forms a near-determinate system distorts the law as much as the view that there is no regularity, no rationality in the operation of justice within the structures of a legal system. The influence of the Realist on the Formalist has been on the whole beneficial since it has served to remind those who are infatuated with the model of a self-consistent and logically 'tight' system that that model does not adequately or realistically capture the inherent

nature of the law. The influence in the opposite direction has also been beneficial: the Realist tends to concentrate on sociological or psychological facts about people who work with the law and overlooks the significance of the logic of legal concepts and the rationality of the reasoning which lawyers and judges engage in. There are, of course, more reasonable versions of each school of thought. There are, as well, intermediate positions which have been expressed. For example, Dennis Lloyd has steered a course between the Formalist and the Realist on this question by rejecting the Realist's view that there is no internal structure to the law, while insisting that the Formalist sees a rigid structure where there is in fact constant movement and adjustment. Hence,

The law is not just a static collection of ascertainable rules by means of which we can at any given moment analyse all the legal implications and relationships which a given concept may entail. On the contrary, the law is a great complex of rules, precepts, standards, and principles in a process of continuous though slow-moving flux. (298-299.)

I shall mention from time to time more moderate views for purposes of contrast, but for the present I want to concentrate on the extremes.

By far the most interesting issue which has exercised jurists--an issue which is a testing-ground for the various schools--has been that of reconciling two countervailing aspects of the common law. We have in the common law the following situation: On the one hand we have the guiding principle of the common law, stare decisis, which requires a judge to preserve the continuity of the law established. The binding force of a precedent is expressed as its ratio decidendi. The ratio is the statement of law, the rule of law, applied to the legal issue raised by the legally relevant or material facts of the case upon which the judgement was based; in a word, the ratio is the legal reason for the decision. The principle of stare decisis and the notion of a ratio of

a case together provide for the continuity of the law. Yet, on the other hand, the common law is engaged in a process of 'continuous though slow-moving flux', a process of extension and enrichment in the light of new cases. Indeed, decisions on those new cases provide the content of, by giving concreteness to, the statutes and precedents. It is, moreover, undeniable that the common law grows and changes with the times, frequently developing along lines which the original legislators did not anticipate, as well as reflecting changes in moral belief and responding to technological, political and economic developments.

The problem then is how we are to understand the common law as providing a legal continuity in the face of continual innovation in the form of new cases. It is easy to see that this is just another form of the question which we have been interested in, viz., How do we reconcile the two intuitive requirements of reasoning, that it be legitimate (by preserving continuity with our present understanding) and that it afford epistemic advance (by providing new understanding)? Suppose we consider how the two extreme positions outlined above approach this question as it is manifested in the more general issue of the nature of the common law.

The Realist will typically seek to undercut the claim that there is a genuine continuity in the common law. Continuity is only apparent, he will argue, because the law does not progress in a logical and orderly fashion. Realists often trace the ancestry of their view on this issue to the Nineteenth Century American jurist O. W. Holmes, and in particular to his opening remarks in his The Common Law:

The life of the law has not been logic: it has been experience. The felt necessities of the time, the prevalent moral and political theories, intuitions of public policy, avowed or unconscious, even the prejudices the judges share with their fellow-men, have had a good deal more to do than the syllogism in determining the rules by which men should be governed. (Holmes, 1-2.)

Holmes said a great deal more on the subject, but the motto that the life of the law has not been logic, but experience, has been developed by modern Realists into the view that an examination of legal reasoning must be an inquiry into the various extra-legal considerations that enter into judicial decisions (vide, J. G. Murphy "Law Logic", 199). There is no point in looking for a continuity-preserving, logical link between precedent case and instant case because the validity of a judicial argument is not a matter of logic at all: a judicial argument is valid just in case the judge was legally exercising his authority in issuing it. There is a semblance of continuity in the law both because there is a more or less constant legal vocabulary and because there is an apparent continuity in the prevalent moral, political and economic practices and policies from which the decisions are ultimately derivable. Still, the argument continues, these practices and policies are undergoing constant change at a fundamental level which becomes apparent if we compare patterns of judicial decision over the course of many generations. The more or less constant legal vocabulary obscures the shifting ground of public opinion. Central legal concepts are employed in old and new statutes and in the binding element of a precedent, the ratio. Yet these concepts are vague, they do not have a fixed meaning which remains unaltered by the changes in public opinion. The law of negligence, for example, is framed in terms of the care which a 'reasonable man' would exercise. But though the judge is in a sense

required to decide that a particular case of negligence is actionable if reasonable care has not been exercised, in fact this is only the weakest of requirements. In the end he must decide what 'reasonable' means in the case before him. Or again, a nuisance is actionable if the thing or action produces in others more discomfort than the ordinary person can put up with. But, the court is not bound to enforce this average standard of comfort and "in the last resort, the application of the rule is a matter of emotion" (G. L. Williams, 579).

The Realist thus voices the complaint that the statutes on the books, as well as the rationes decidendi which are extracted from previous cases, cannot serve as major premisses in a deductive argument. This in itself is an unobjectionable point, indeed I shall be reaching this same conclusion myself in what follows. What is objectionable, however, is the thesis which the Realist goes on to assert, namely that there is no rational basis for the development of the common law. The position at times is reminiscent of the emotivism of A. J. Ayer and C. L. Stevenson, at least with respect to the general scepticism both views express concerning the possibility of objectively deciding the correct judgement to be made in particular cases. For both the Realist and the emotivist justify this scepticism by pointing out how deductive reasoning either fails to get a grip on the problem or else is applicable only if we irrationally posit guiding principles. One of the central complaints of the Realist, for example, is that the notion of ratio decidendi misleadingly suggests that the judge is merely applying unambiguous principles to instant cases, that he is logically bound by precedent. But, the Realist insists, this is not at all the way it works, since rationes are too open to interpretation and controversy to

be able to serve that deductive role. They can be ignored, or in the jargon, distinguished: "...the humblest judicial officer can disregard the most authoritative declaration of the House of Lords unless he considers that the precedent cited is 'on all fours'" (Allen, 47-48).

It should be noted that this particular aspect of the Realist position in an important sense reflects its origins: in the United States, and to a lesser extent in Canada as well, the doctrine of adherence to precedent has always been applied with far less rigidity than it has in England. So to a large extent the American Realists have indeed been reporting 'realistically' the workings of the common law. Nonetheless, in English as in North American law the opportunity to ignore or otherwise reject rationes is a built-in feature of the machinery of the common law. In practice, the refusal to be bound by precedents can take one of the following forms: (1) A judge may distinguish his case from a precedent cited by arguing that the material facts of his case are not sufficiently analogous to the material facts of the precedent. (2) A judge of a Court of Appeals (or higher court) may overrule a precedent by, as it is usually put, 'restating the law correctly'. (3) Judges of the Court of Appeals may cite the per incuriam doctrine (from Young v. Bristol Aeroplane Co., Ltd. (1946) 1 All E.R. 98), claiming that they are not bound to follow a decision on the grounds that it was made 'through lack of care'. And, of course, (4) a judge may reinterpret the ratio of a precedent (vide Cardozo 158 ff. and, for example, Dias 63 ff.).

In any event, the Realist claims that the judge can and should look elsewhere for the bases of his decision. The justice which he brings to his decision is not justice compelled by logic; it is rather a justice which reflects the standards of his society, standards which need not

themselves be rationally justified. As Cardozo has suggested, at some point in the judgement-making process reasoning ends. At this point: "History or custom or social utility or some compelling sentiment of justice or sometimes perhaps a semi-intuitive apprehension of the prevailing spirit of our law must come to the rescue of the anxious judge, and tell him where to go" (43).

It has been variously argued that the extreme Realist in the end just fails to account for the unity and continuity which seems operative in the actual working of the common law. It is a historical point that the entire body of laws relating to tort (which today encompasses various forms of negligence, nuisance, defamation, trespass and misrepresentation, as well as false imprisonment and malicious prosecution) grew out of the single writ of trespass vi et armis et contra pacem domini regis (by force and arms and against the King's peace) originating in the 13th Century. But this argues against the Realist's position. For although it is obviously true that the separate torts of negligence, nuisance and the rest developed as a result of changes in social conditions and an awareness of the variety of ways in which the King's peace can be disturbed, still there was a underlying rationality to this evolution. New crimes do not sprout fully formed from the minds of men, they are extensions and modifications (obviously in the light of changing conditions) of old crimes. Similarly, new decisions which help to extend and modify the existing conceptions of negligence, nuisance and the rest fit into the fabric of the law. For surely, if a particular judicial decision did not bear any relation to previous decisions in analogous cases, then there would be clear grounds for appeal. There is, then, need for an account of the systematic structure of the common

law, and it is this which the Formalist was trying to exhibit all along.

But the goal of Formalism has always been that of fitting legal systems into a deductive model. Consequently, the Formalist sees his task as that of analysing legal concepts and propositions in a manner which will make them more determinate. The idea is to force the common law in the direction of a closed determinate system. This is of course an ideal, and most Formalists recognise that. They are nonetheless caught in the grips of what R. Pound has called the aspiration of 'mechanical jurisprudence', in goal of generating every judicial decision deductively from statutes and precedents. There is little doubt that if the law could be mechanised in this way it would enjoy that property held in common by closed determinate systems--self-consistency would be lost by incompleteness, for by calcifying the law one is destroying its capacity to respond to changes in social conditions and other relevant features of new cases. Take for example the vast range of novel cases which are brought to the law whose novelty is the result of technological advance. When the radio became an effective mode of communication there was a flood of cases which hit the courts: Is it libellous to say rude things about your neighbour on the radio? Is it trespass to broadcast without permission the scores of cricket games being played on a private field where an entrance fee is charged? Is it gross negligence to broadcast that some fictitious catastrophe has occurred, like the landing of martians in upper New York state? These cases involved technological considerations which could not have been dreamed of by the original framers of libel, trespass and negligence laws (some of which were several hundred years old). But technological change is just the obvious source of novelty. Equally important are changes in social

awareness of injustices previously ignored, changes in economic or political conditions, and so on. The Formalist will respond that the novel cases either require new legislation or else fall unambiguously under existent statutes. It is true that the statutes are constantly being reviewed and revised and that new statutes are being added. But although there are actual changes in the statutes, the common law functions in such a way as to accommodate, without the need for explicit legislation, new and novel cases. The source of this accommodation is judicial 'creativity', the interpretation and extension of existing law by judicial decision. The Formalist is typically very harsh in his rejection of the need for judicial creativity and insists that a single ratio decidendi can be extracted from each precedent case which either has or does not have a clear application in each subsequent case.

Now, much of the Formalist's argument concerning the nature of the common law rests on this last claim. If the Formalist is to preserve the model he favours, there not only must be a single ratio decidendi for each precedent case, but that ratio must also be able to serve the function of a major premiss in a deductive argument, the conclusion of which is the legally correct decision to be made in the instant case. Judicial reasoning must take stare decisis seriously, the argument goes, the reasoning must lead to a decision which is 'principled', that is "based on reason, not fiat, on doctrine impartially applied, not on ad hoc intuition or whim" (Henkin, 302). Few would want to make the claim that judicial decisions ought to be made on the basis of fiat, ad hoc intuition or whim (even the Realists, toward whom this pointed remark was directed, speak of socially responsible, and experience-trained choice). Nonetheless, the particular logical role

which the Formalist wants the ratio decidendi to serve is simply out of reach. The problem is with the notion of the ratio decidendi itself, although there is a problem with the notion only if one assumes that it must serve the role of a major premiss in a deductive argument.

The ratio decidendi of a case is the reason which is binding on future cases. The ratio is binding on future cases because it was deemed necessary for, or the basis of, the precedent decision. The ratio is thus distinguished from other obiter dicta which are auxiliary reasons used for the judgements which are not binding on future cases. To determine the ratio, on one popular view (vide, Goodhart), one must identify the facts which were in the precedent court's view 'material' or relevant and then find the reason which explains the decision made in terms of these facts. What one is after is a reason for treating a precedent and an instant case as relevantly similar, sufficiently similar, that is, to warrant the citation of the decision made in the precedent case as a reason for deciding in a similar fashion in the instant case. The form of argument employed here is argument from analogy (although stare decisis requires the analogy to have legal force in one direction only: the court is required to apply past decisions to instant cases, standing by the past decisions). In practice, it should be noted, this form of reasoning is not treated as if it were conclusive or compelling. The court is not solving a problem or making a calculation when it considers competing precedents in an instant case; it is rather making a judgement, the judgement that there is a relevant similarity between the two cases. Part of what is open for reasoning and judgement is what the reason for the decision in the precedent case was--and this even when a reason for the decision was

explicitly stated by the precedent court. A precedent is a particular case, a unique collection of features and circumstances--some of these features and circumstances will be material to the case and others not. The instant case is also a particular case. But if the court is not making decisions about abstract cases of 'pure' negligence--cases shorn of all the particular details which, if brought together in different combinations may cast doubt on the certainty of judging one way rather than another--then it is not making decisions about ultimately unique cases either. Not all of the particular details of a case are material. And that there is a distinction to be made between material and non-material facts makes it possible for analogies to be drawn. A ratio is essentially a guide for spotting the relevant facts, or more correctly, the relevant sorts of facts involved in past cases.

In the hands of the Formalist, however, ratio decidendi becomes an extremely problematic notion. The Formalist conception of the common law requires that the ratio include an explicit statement of all and only the material facts of the case. What the Formalist is after is a sort of formula akin to a scientific law. The judge is supposed to have a list of rules at his disposal and his job is to apply one (and only one) of these rules to the instant case. But the ratio does not, and could not, act like a formula for deciding cases, nor is it designed to. But if the Formalist persists in requiring the ratio to act like a formula for deciding cases, it becomes very difficult to explain how the ratio can be constructed. For one not only has to isolate all and only the material facts of the precedent cases, one also must decide at which level of generality these relevant facts should be preserved.

The difficulties involved in constructing an explicit ratio can

be seen more clearly if we look at a case which is frequently mentioned in this context, the case of Donoghue v. Stevenson (1932) A.C. 562; cf. J. Stone, 269-271, and R. Cross, 182-186). In this case a Scottish widow one day in 1932 was supplied by her local shopkeeper with a bottle of ginger beer, part of which she drank and the rest of which she poured into a friend's glass. As she poured the ginger beer, however, a decomposed snail floated out of the bottle. As a consequence of having drunk part of the contaminated contents of the bottle, the woman alleged that she contracted severe gastric pain. The bottle was made of opaque glass so that the condition of its contents could not have been ascertained by inspection. The question before the court was, given that the shopkeeper was merely the retailer, whether a manufacturer owed a duty of care to the ultimate consumer. The court ruled that it does. The ratio of this case is usually given as follows:

A manufacturer of products, which he sells in such a form as to show that he intends them to reach the ultimate consumer in the form in which they left him with no reasonable possibility of intermediate examination and with the knowledge that the absence of reasonable care in the preparation or putting up of the products will result in an injury to the consumer's life or property, owes a duty to the consumer to take that reasonable care.

Now, first off, it is clear that the particular facts of this case do not form part of the ratio. The reason for the decision does not hang on this Scottish widow developing this sort of malady from this sort of ginger beer in 1932. But the history of the cases which have followed Donoghue v. Stevenson shows that various levels of generality for the various material facts of the case were found to be appropriate. For example, the agent of harm in the precedent case was decomposed snail in an opaque bottle of ginger beer. Yet, the reason the precedent

court had for ruling against the manufacturer could equally be applied (and has been applied) to cases where there was another sort of noxious foreign substance in another sort of product sold for public consumption. Thus in Grant v. The Australian Knitting Mills ((1936), A.C. 85) it was ruled that as the plaintiff had contracted dermatitis from wearing pants containing an excess of sulphites, the presence of which was due to negligence on the part of the defendant, the defendant owed a duty to care to his ultimate customer. In Haseldine v. Daw ((1941), 2 K.B. 343), further analogical extension of the original ratio was affected. Here the plaintiff sustained injuries in a consequence of the collapse of an elevator in which he was ascending on his way up to an apartment of one of the tenants of a large building. It was ruled, once again on the basis of the Donaghue v. Stevenson decision, that a repairer of an elevator stands in the same position as a manufacturer on the issue of duty to care to ultimate consumers/users. Other directions and other vectors of analogy have been followed: in the precedent case the noxious element was not discovered by the retailer because of the opaque bottles; but the reason for the original decision has been ruled to apply to cases where the noxious element was not discovered by anyone who could reasonably be expected by the defendant to inspect the product. The permutations of material facts which can be, and have been, judged relevantly similar, along different vectors of analogy, is obviously quite extensive.

But the point here is that the Formalist's insistence that there is only one, unambiguous ratio of a precedent must be rejected: in fact, the ratio serves as a starting point, but not as a major premiss in some judicial syllogism:

...there is no fixed ratio decidendi; it is not only the ruling given by the deciding judge for his decision, but any one of a series of rulings as elucidated by subsequent interpretation. The pronouncement of the judge who decided the case is a necessary first step towards ascertaining the ratio, but the process by no means ends there; the subsequent interpretation is at least as significant, and sometimes more so. (Dias, 63-64).

A court ruling that a particular action or omission constitutes negligence is providing future courts with an example of what it means for someone not to have exercised 'reasonable care'. Citing the case as a precedent, a future court is acknowledging that, if the instant case is set alongside the precedent case, legally significant features of the instant case are highlighted. When these features are brought out in this manner, one is not supposed to be able to detect an identity relation between the two sets of material features; indeed, for most cases, the two sets of features will not in any straightforward sense resemble each other. An automobile in bad repair can be a noxious physical object, but no one would venture the claim that it resembles an opaque bottle containing decomposed snail (cf. Oscar Chess v. Williams (1957), 1 W.C.R. 370 (C.A.)). Yet, reasoning proceeds on the basis of a legally significant analogy between the two sets of material features, those of the precedent case and those of the instant case. And it is this analogy which makes reasoning from precedents possible. The Formalist will have nothing to do with talk of analogies, he is after deductive relations. But because of the nature of the field of the common law, strict deductive relations are not always forthcoming. As H. L. A. Hart has noted with respect to negligence,

...owing to the immense variety of possible cases where care is called for, we cannot ab initio foresee what combinations of circumstances will arise nor foresee what interests will have to be sacrificed or to what extent, if precaution against harm is to be taken. (129.)

But, significantly, the Realist is operating under the same assumption as the Formalist. The difference is that whereas the Formalist thinks he sees the possibility of determinate principles in the common law, and therefore the possibility for genuine reasoning on the judge's part, the Realist denies that such principles exist, and therefore that no genuine reasoning in the common law is possible.

I have already mentioned that several legal philosophers have tried to avoid both of the extremes of Formalism and Realism in their characterisation of legal systems and legal reasoning. I want to consider some of the more specific recommendations of these jurists in subsequent chapters; but here I want to isolate some of the dominant themes of this moderate approach to the question of reconciling continuity with change in the common law.

Briefly, the moderate view is that we do not have in the law determinate principles, either in the form of rationes decidendi or in the form of statutes. The instances of rationes and statutes are not uniquely recoverable, i.e., the relationship between rationes and statutes and particular cases of each is not the relationship of principle and instance of principle. Decisions about particular cases do not, therefore, follow merely by deduction from the laws, whether legislator- or judge-made. There are no truly determinate principles in the law for the reason that the key concepts of the law are themselves non-determinate. We cannot decide in advance the necessary and sufficient conditions for an act or omission being negligent, exhibiting a failure to exercise reasonable care, being cruel, libellous, or unfair. Nor are there explicit definitions for concepts which are centrally important in particular kinds of cases, concepts like just

cause or excuse, undue hardship, defamation, mens rea, insanity, contract, wrongful arrest, fraud, possession, trespass, duty to care and a host of others (most of which are not presumed to have a technical, legal meaning). On the other hand, a judge ought not to think that he is giving a reasoned decision unless he views the instant case in the light of a full range of similar cases and, of course, the relevant statutes. The concepts employed in the law may be non-determinate, but this does not mean that no reasoning about these concepts is possible. There are cases where reasonable care is clearly exercised, and there are cases where it is clearly not exercised. The endpoints of the spectrum of cases of exercising care are clear enough; the problematic cases are those which fall somewhere in the middle of the spectrum. To decide the case before him, the judge must view the instant case in the light of the law, i.e., in the light of previously decided cases which provide the logical links between the relevant legal concepts. The basic view, then, is an expansion of Felix Frankfurter's famous admonition: "The particular in isolation is meaningless; the generalisation without concreteness, sterile" (quoted by H. Jones, 427).

Now if, as the moderate view contends, there cannot be determinate principles and concepts in the law, then deductive reasoning from principles cannot constitute the whole of legal reasoning--although it has a role to play at various stages of the deliberation. The central problem for which judicial reasoning is required is whether the facts of a particular case warrant the application of certain legal concepts to the case, whether it is just to treat different cases as though they were, from the point of view of the law, the same. To be sure, it follows deductively that if all cases of failure to exercise reasonable

care are actionable as negligence, and if this case is one where there was a failure to exercise reasonable care, that this case is so actionable. But this bit of reasoning is possible only after all the real work of judicial reasoning is over (cf. Cross, 180 f.). If we have reason to say that this is a case of failure to exercise reasonable care, then the judicial question has been solved, the decision has been made. The statutes concerning negligence provide the conceptual, and so the logical connections between negligence and reasonable care; precedents concerning negligent acts provide a cluster of cases which exemplify and make concrete this conceptual connection. But neither statutes nor precedents provide formulas for determining how that connection will be made concrete in future, as yet unheard, cases.

Another central theme of this moderate view is that we must reject the suggestion that, in the absence of conclusive and compelling reasoning from principles, the judge must resort to intuition or attempt in his decisions to do no more than reflect the predominant views of the society to which he belongs. If a court rules that a statute imposing liability for negligent actions should be extended to apply to negligent statements, it will not be a ratio of this decision that 'the time is right' for such an extension of the concept of negligence, or that a majority of people in this society would be willing to accept such an extension. Social needs, conceptions of human welfare, and economic conditions obviously change over the years; and these changes are casually connected to changes in the law. And, furthermore, these changes may explain the need for expanding the law (up until very recently, for example, international law made no provision for questions concerning the acquisition of sovereignty over 'territory' on the moon).

All this can be granted. But the reasoning advanced in support of a decision must be distinguished from an explanation of why it became necessary to make a decision on the case. A judicial decision is reasoned when that decision can be shown to fit into a pattern of similar cases, to accord with the conceptual connections embodied in statute and precedent. But the very content of both statute and precedent derives from a past history of cases. In the law, as in every field of human knowledge, the advance which bears no relation to previous cases, the question which is so novel as to stand completely outside of the recognised sphere of physics, morality or the law, is neither an advance nor a novelty; it is a mystery.

Legal generalisations contribute to legal reasoning only in so far as they can be given concreteness, and particular cases can only be reasonably adjudicated if they are viewed against the background of a network of other cases. Continuity in the law is provided by the structure of statutes and precedents, and in particular by the analogical connections which link cases together into a network which gives concreteness, and content, to the statutes. Change in the law is provided by the elasticity of the structure, and in particular by the non-determinacy of the central legal concepts. Reconciliation of these two features of the common law thus rests on a non-hierarchical and essentially non-deductive conception of the system of the law.

One last element of the picture of the common law here being developed should be noted. Since a legal system is not (in our sense) a closed system, there are bound to be gaps in the fabric, that is, legal problem-situations for which no authority, statute or precedent, exists (casus omissus). There is in modern jurisprudence great support

for the prohibition of what is called non liquet, that is a rejection of the claim that a court (particularly in matters of international law) can refuse to decide a case on the ground that there is no precise authority applicable. (French Civil law makes this a matter of law, requiring a judge to decide a case even where the law is silent or obscure on the point at issue.) This theory of the 'logical plenitude of the law' (so-called by G. W. Paton, 149 ff.), is motivated by the belief--well supported by the historical progress of the common law--that the law has an inherent ability to grow and adapt to new conditions and circumstances.

Now in the light of the continuity/change problem we have been considering, the view that the law is a logical plenum underscores the viability of, and necessity for, a case-oriented conception of legal reasoning. For when no statute or precedent is directly applicable to a particular case, the judge must search for, and is typically presented with, competing precedents which are in the logical 'neighbourhood' of the instant case. It has been argued by O. W. Holmes that nearly all of the cases in common law fit into the gaps and spaces of the fabric of the law, and hence that legal reasoning is inherently 'interstitial'. Yet, enough has been said to support the point that even when a particular case seems to be covered by a statute or precedent, the authority thus provided is not logically irrevocable. The operation of the court is the operation of a forum, and competing precedents are brought up in order to bring different features of the instant case into light. Statutes are not formulas, and no case will be identical to a precedent. Still, statutes and precedents are laws, and cases must be decided. But the sort of reasoning called for involves competing

pictures of what the instant case is like. Features of the case are isolated and viewed as material, and these are brought together to produce a complete picture of what the instant case is a case of. At each step a reason is given for the composite picture and these reasons, as Wisdom has suggested, "are like the legs of a chair, not the links of a chain" ((1), 203).

Now all of these considerations point to a conception of legal reasoning which not only rejects the view that such reasoning is essentially reasoning from principles, but also questions the assumption that it is deductive or inductive. Julius Stone, for example, warns that

...the limits of formal logical reasoning in appellate decision-making, cannot and should not be, and never has been the limits of all reasoning there. (333)

Primarily in light of the realisation that the ratio decidendi of a case is often but a guide for spotting relevant fact-types, positive accounts of the nature of judicial reasoning have been offered which emphasise the central role of reasoning from particular cases. Several jurists have expressed a willingness to agree with Edward Levi's claim that "the basic pattern of legal reasoning is reasoning by example. It is reasoning from case to case" (1). In short, the view has been expressed that 'mother's method' is indeed the key to understanding the reasoning which lawyers and judges use, and use successfully. The view has developed out of a rejection of the two extreme positions we considered above, and has steered a course between the one extreme of succumbing to Euclid's disease and the other of resigning oneself to a picture of the law as an incoherent collection of statutes and decisions. A closer look at the actual subject-matter of legal reasoning, a sensitivity to the working methodology of the courts, and an avoidance of the

'craving for generality', has lead to a view of reasoning which emphasises the importance of reviewing particular cases, of reasoning from analogies and not from principles. I want now to leave the theoretical, jurisprudential sphere and develop some of the analogies which obtain between legal reasoning and reasoning about other matters.

Questions about Concepts

Our excursion into jurisprudence in the last section uncovered reasons for thinking that judicial reasoning cannot be adequately characterised as reasoning from principles in a deductive fashion. These reasons had to do with the sort of system a legal system is, with the non-determinate nature of the statutes and rationes decidendi and the open-textured nature of the concepts which are employed in the law. Still, if the evidence seems to argue against the Formalist's dream of a 'mechanised' jurisprudence, it also argues against the Realist's scepticism about the possibility of genuine justification and proof in the law. For the questions which are argued in the courtroom, the questions which are there to be decided, are questions which require reasoning. There is fact-gathering, of course, but there is more--there is reflection over the facts. The courtroom is a forum which serves a social as well as a rational function. The forum protects the interests of the parties involved and the community by making sure that competing arguments and competing interpretations of the facts are before the court. But the forum also insures that the issue is debated, that reasons for and against guilt or innocence are aired and criticised.

But we also saw in the last section that the reasoning which lawyers and judges engage in is not exclusively deductive in form. It was suggested that the important reasoning, the reasoning which carries

the court towards a decision, is reasoning from case to case, reasoning by comparison and contrast. It is reasoning aimed at the question whether a certain act or omission is of a certain, legally significant sort. I suggested that the fact that the reasoning which is employed in a court of law is not exclusively deductive is explainable in terms of the non-determinate nature of the legal 'field'. But also important here is the nature of the question which is being discussed, the object and issue of the reasoning. A closer look at the sort of question which requires judicial reflection will reveal the parallels which exist between some legal questions and most of the questions which arise in non-legal contexts which cannot be answered by an appeal to a principle, rule or standard. Such questions are noted for the perplexity they cause, a perplexity which results not from the absence of information but from an initial ability to put all of the pieces together--questions, in short, which call for reflection. Suppose, then, we turn to the law of torts again.

It is doubtless the case that every form of activity or inactivity carries with it the possibility of harm to others. But liability for fault in the case of negligence involves the idea that the risks created by an act or omission be unreasonable. One factor which has developed into a sort of test relevant to the question of the reasonableness of risk is that of the foreseeability of possible harm. In several cases the question of whether the defendant could have, and so should have foreseen that his act or omission would result in harm to others has been one of the principal questions at issue during the trial. And the question is a straightforwardly conceptual one, it is a question about what it means to be able to foresee possible harm.³ As the test of

foreseeability is used again and again, the legal content of the concept is delimited, not by means of exceptionless necessary and sufficient conditions for the possibility of harm being foreseeable, but by means of individual manifestations of the correct or incorrect application of the concept in concrete and fictional cases.

In the case of Law v. Visser ((1961) Queensland R. 46), for example, the defendant was a motorist who, while driving at night and speeding, saw in the road ahead a large object which looked to him like an abandoned bundle, or possibly an animal killed by a previous car. The defendant made no effort to avoid the object and ran over it. The object proved to be the plaintiff, asleep in the road in a drunken stupor. The court asked itself whether the defendant ought to have foreseen the possibility that the object in the road was more valuable than he thought, or even a human being, however unlikely that might be. This question naturally lead to another conceptual question: How much care should the reasonably careful driver in such a circumstance exercise? In another case, the defendant was a truck driver who had loaded his truck in such a way that as he passed under an overhead bridge a container was knocked off and fell on the plaintiff. The plaintiff, shortly before the accident, had been on the truck as a trespasser and at the time of the injury was running slightly behind it attempting to steal a second ride (Farrugia v. Gt. Western Ry. (1947) 2 All E.R. 565). Should the truck driver have foreseen the possibility that someone would be running directly behind his truck when he took the chance and tried to get under the bridge without knocking anything off? Involved here as well is the counter-vailing consideration of trespass and the responsibilities which a driver has to the safety of possible trespassers.

Rather than dwell on the details of the courts' reasoning and ruling in these two cases, I want to consider the nature of the questions which were at issue. In a sense what I want to do is to demystify them, to show how very like other, non-legal and ordinary, questions they are, despite the fact that they are found in a legal context and were argued there. Two points can be made in that direction: (i) First of all, in both cases the central issue involves a question of meaning. The test of foreseeability as applied in these cases is not some exotic legal apparatus which works in ways which only lawyers or judges could possibly understand. It is rather the process of seeing whether a concept can be correctly applied to a set of circumstances. No special, legal sense is assumed for 'foreseeing the possibility of harm'; the phrase is not like, say, the terms 'person' or 'larceny' (or the distinction between malfeasance and non-feasance) which are part of a specialised, legal vocabulary. (Furthermore, the specialised legal vocabulary is not very far removed from ordinary language anyway, consisting as it does of terms used to draw distinctions which are often made in ordinary discourse.) Rather, the judge on reflecting over the case of the motorist is reflecting on our ordinary notion of foreseeability, although his concern is over a particular legal case. He does not rely on a precise formula or a definition. He is seeking definiteness, of course. He is required to judge one way or another. And as a judge he is interested in a legal 'ought', not a moral or a prudential 'ought'. But his reasoning is directed to the conceptual question of whether the motorist was negligent when he didn't think twice about running over the bundle in the road, whether the truck driver was negligent when he took the chances he did.

Some legal cases hang on questions about the correct interpretation of very ordinary words indeed. An interesting example of this is a recent appeal case in which the issue was an insurance company's refusal to pay a claim on a policy insuring against 'storm, tempest or flood' (Young v. Sun Alliance and London Insurance Ltd., The Times Law Report: 15.5.76). The issue here was over what anyone would agree is a quite ordinary notion, the notion of a flood. The claim was made on the basis of water damage done to a house which was caused by the ingress of water into the basement from an underground watercourse. The appeal was lodged against a previous judge's interpretation of the word 'flood' which was not, as it was argued, wide enough to cover the facts of the case. It was argued that any householder who could not walk about in a room in his house without being in three inches of water would certainly be entitled to describe the state of affairs as a flood. The appeal was dismissed on the ground that the sort of phenomena which was being insured against necessarily involved an element of 'violence and suddenness' and seepage of water was neither violent nor sudden. Lord Justice Cairns, in concurring, pointedly remarked that the "way to interpret the word 'flood' was to ask oneself what meaning the ordinary intelligent man would give to it." This appeal to the ordinary intelligent man is significant because it exemplifies the sort of reasoning which typically takes place in the law, namely reasoning about the meanings of words from the perspective of ordinary use.

(ii) Because the issue is a conceptual one, the judge's reflection takes place after the stories have been told and the evidence is in. The judge must reflect over the facts of the case as they have been laid out during the trial. He is not attempting to uncover another piece of

evidence about the case which did not come out before. He is not, in other words, a detective; he is not investigating the alleged facts of the case, checking stories and looking behind the scenes. All of that he presumes has been taken care of before he begins his deliberation. All that the judge deliberates over is in front of him; he has to make sense of the pieces. But of course the same applies to the rest of us when we consider similar sorts of questions. Naturally enough, the judge's role, the part he plays in the institution of the law, makes for some differences between his deliberation and ours. For, as we saw in the last section, he must make a decision one way or another, whereas we can plead that not enough evidence is at hand to decide. In the law as in morality there are bound to be cases where the correct thing to say is that the question of right or wrong, guilty or not guilty is unanswerable. Given the complexity of moral and legal questions, it is reasonable to suppose that for some issues the correct answer is, There is no 'yes or no' answer. That legal issues cannot be left to dangle in this fashion, that particular cases, despite their complexity, must be decided one way or another, makes for a crucial difference between the law and morality. Another difference, also significant, is that the judge's decision is law and forms part of the interlocking network of decisions which are cited or reviewed: the law is both a system and an institution.

Moreover, the judge's reasoning is at times limited by constraints which are imposed on the kind of issue in question, constraints which are legal in nature, and constraints which the rest of us, outside the legal context, do not feel. Although the evidence is there for all to see, the judge is an interpreter of the law, the law as it is. Nonetheless,

new evidence may surface which may require a new trial. If the evidence is both new and relevant then the pieces of the puzzle stand a good chance of forming a different picture. And, significantly, if a new trial is ordered, the case is treated anew; for it is a different case.

The typical question which requires judicial reasoning, then, is a question about the applicability of a legally relevant concept which is only significantly asked after all the evidence is in. Such a question may join with others which require reflection and the answers to these questions jointly lead to a decision. There are non-everyday aspects to judicial reasoning since it takes place in a very specific environment and is subject to strict procedural rules. For example, the common law is suffused with so-called 'rebuttable presumptions of fact' which serve to fill in gaps in the factual framework of a case. These presumptions, recurring instances of the operation of circumstantial evidence, serve to supply (unless rebutted) facts which, in a non-legal context, we might be unwilling to assume. Thus it is a presumption of fact that if a person is in possession of recently stolen goods, he has guilty knowledge of the source of those goods. This is a legally sanctioned inference which, although reasonable enough, we are not, outside of the legal context, required to make. The precedent system, moreover, serves to codify, in a manner which is not done in ordinary contexts, the decisions about previous cases which together give meaning to legal and legally relevant concepts.

This last fact goes a long way towards making sense of the important notion of ratio decidendi: every decision helps to give life to the concepts of carelessness, risk, harm, reasonableness and the other key concepts which figure in the common law. The legally relevant meanings of these concepts grow and expand with the accretion of decisions about

particular cases. There is, moreover, continual refinement as well as expansion: judicial decisions are often corrected in the light of later cases (this less frequently occurs, at least overtly, in English law). The legal forum is not confined to the courtroom; it covers the whole history of the law. Cases are brought forward as precedents and compete among themselves--in the light of the particular features of the instant case. Rationes constitute the links between cases, threads of continuity.

Now it might be objected that deduction provides the overall structure of a judicial argument, and, moreover, rationes, statutes or both count as major premisses in the resulting deductive argument. I certainly would not want to deny that judicial decisions, the completed argument, can be expressed in terms of a deductive argument, or that, once put in that form, rationes and statutes play the role of major premisses. But what I would like to claim is that these deductive arguments can be constructed only after the real work of interpreting rationes and applying statutes in the light of the particular features of instant cases has been completed. Before it becomes possible to utilise deductive reasoning, the rationes in particular can only serve the role of 'for the most part' guidelines, guidelines based on what in the precedent court's eyes were the features of that case which ought to have on-going legal significance. At this stage in judicial reasoning, reflection is called for. The reasoning required is reasoning directed to, for example, a claim like 'Features in this case, analogous to those of this precedent tend to support the view (are reasons in favour of saying) that this is a case of negligence'.

It is tempting when considering the legal forum to point to procedural rules, the directions and other formal features of the

institution of the law and to use these to support the claim that there is something exotic, or at least extra-ordinary, about the reasoning which leads to legal decisions. There are quite explicit rules, there is no doubt. In criminal proceedings the prosecuting counsel always begins first, and he is followed by prosecution witnesses who are subject to examination-in-chief by prosecuting counsel, cross-examination by defending counsel and re-examination by prosecuting counsel. There are also explicit rules for pleading, for formal admission and for the determination of the onus of proof. Especially in criminal law there are certain formal presumptions made with regard to the determination of where the general burden of proof shall lie (e.g., the presumption of innocence) and irrebuttable presumptions concerning the nature of the crime and the status of the criminal (e.g., it is presumed that no child under the age of ten can be guilty of any offence). And lastly, there are explicit rules governing the force and admissibility of the various forms of evidence. All of these features of the working of the court contribute to the plausibility of the view that legal reasoning is substantially different from reasoning outside of the law.

But it is clear that procedural rules and presumptions are intended to guarantee a fair hearing and to preserve a continuity of practice. These rules and presumptions delimit the boundaries of legally valid procedure; they determine the course of the argument, but not the nature of the reasoning which is required. For it is significant to observe that the sorts of questions which we are presently discussing are questions which are put in the hands of juries. Their reasoning is neither exotic nor extraordinary, although the directions of the court severely limit the range of considerations which they can treat as

relevant. It is the duty of the judge to specify the question over which the jury is to deliberate, or indeed, to direct the jury under certain circumstances to reach one verdict rather than another. The jury's job is not to determine the merits of the law, nor can they speculate on questions of procedure. But in putting the question to the jury, when he is not explicitly directing that they reach a certain verdict, the judge is not thereby informing them of some set of legally specified, necessary and sufficient conditions for, say, foreseeability, carelessness or reasonableness. Rather, the question put to the jury is one of fitting the facts as presented during the proceedings into a picture which warrants calling the case one of negligence, or which does not warrant this.

This role of the jury in the legal process is thus of great significance to the question of the nature of reasoning in the law. For it becomes clear that the reasoning which is expected of the members of the jury, as reasonable men, is not reasoning which is carried out against the background of a legal system, or even against the background of that part of the law which deals with the tort or offense in question. Jurors are not versed in the law; indeed, lawyers are specifically excluded from jury duty. They are presumed to be, and it is enough that they should be, reasonable men. The questions which are put to juries are questions which require reflection over the evidence which they have heard; and the question is typically one of the applicability of legally relevant concepts: Is this case one of care or the lack of it, of the reasonable exercising of responsibility or not, of harm which ought to have been foreseen or not? These are questions of fact which arise after all the facts that can be gathered

have been gathered. And it is the jury's duty to construct out of the answers to these questions a verdict.

Some of the questions which juries are asked to consider are extremely difficult to answer one way rather than another. The case may involve a borderline situation where good reasons can be given for saying that, say, it was a case of carelessness and good reasons can be given for saying that one could not have expected more care from any reasonably careful person. And some of the questions which are not given to juries are not because they involve issues which, in the opinion of the judge, could be the source of legal controversy of a sort ordinary people could not be expected to be sensitive to. Not every case is a leading case, of course; nor is every question a question about a borderline situation. Still, every case requires reflection. Hence the standard of judicial proof--proof beyond reasonable doubt. To determine guilt or innocence beyond reasonable doubt is not to determine it beyond all reason for doubting; and reflection is needed to determine when a doubt is reasonable.⁴

I am not suggesting that a question like, Were the defendant's actions, as described in Law v. Visser, the actions of someone exercising reasonable care?, is a question which can be reflected upon and argued about in a legal context in ignorance of the law of torts. For, in the context of the law the question is typically asked in terms of negligence in the sense of other cases of negligence found in law reports, not in terms of negligence sans phrase. The judge's instructions to the jury are designed to set out whatever legal sense of a concept like negligence is relevant. The question, the conceptual issue, is then in the hands of the jury. Still, outside of a legal context, we

can reason about the status of the defendant's actions; and the reasoning we would engage in would be parallel to, although not bound by the same constraints as, the judge's, or the instructed jurors', deliberation. Another point is important here. Given the nature of the question being considered, whether the case actually occurred or not is quite irrelevant to the manner in which we would proceed. Doubtless the discovery that one of the cases found in a standard text on leading cases in the law of torts was pure fiction would cause a stir in legal circles, and great consternation among judges who have employed the case. But pure fiction or not, the case as described gives us enough at least to begin a rational consideration of the conceptual question at issue. The point being made here is just the point made in Chapter II where it was argued, in the case of deductive inferences, that the falsity or fictitiousness of the premisses is neither a good nor a sufficient reason for thinking that the argument which followed is a bad argument. This point will become increasingly important in what follows.

The reasoning which is carried out by a judge or a jury over questions about the applicability of non-determinate concepts like foreseeability, carelessness, reasonableness, risk, harm and the rest is perfectly respectable reasoning; but it is not deductively conclusive reasoning. But the kind of reasoning which is carried out in the courtroom or the judge's chambers is reasoning which is particularly suitable to the kinds of questions we have been discussing. What has to be taken seriously is the existence of questions which require reflection, but are such that the reflection is conceptual, both a priori and non-necessary. In the law these questions are the central objects of debate and deliberation, they are the key ingredients of judicial decisions. But are there

other, non-legal, examples of these sorts of questions?

Consider another sort of case: Suppose I have known X for many years. I have watched him react to various situations and I am familiar enough with his mannerisms, his style, to be able to mimic him successfully in the presence of mutual friends. Yet, despite my knowledge of his ways, if I am asked to describe his personality and the way he carries himself generally, I may very well have the problem of trying to capture what it is about him which makes him X: is he dignified, or is that pompousness; is he meddlesome or sincerely concerned with the problems of others; is he a bit of a swell or just self-conscious about his stylish looks; is his tendency to find fault in others frankness or newsmongery; are his occasional falsehoods conscious fabrications or careless exaggerations? If these questions are put to me I may watch him more carefully to see if I can catch any clues that I may have overlooked in all these years of acquaintance. But after I have observed him I may still be perplexed: the facts are before me but I still can't put them together in a way which will count one way or another. My perplexity may vanish in a flash, however. I may catch X practising stern and dignified looks in a mirror and realise that there was all along something not quite genuine about his stateliness. My problem is one of adequately describing X, and I attempt to capture the man's personality traits by trying on concepts like dignified or pompous to see if they fit the facts. I compare his actions and looks with clearer examples of dignified looks or pompous actions to see similarities or differences.

The problem of seeing a person's ways in the proper light is common enough in our everyday dealings with people. But the problem is

analogous to the difficulty of adequately capturing the personalities and motives of fictional characters. The literary critic must be a good judge of character. Graham Greene's 'confidential agent', for example, is an intriguing mixture of the cynical and the defenceless. He is prone to muse cynically that

...in a happy life the final disillusionment with human nature coincided with death. Nowadays they seemed to have a whole lifetime to get through somehow after it....⁵

And yet his dealings with others betrays a sort of vulnerable optimism. We can pour over the text looking for clues, and we may find them. But to get a complete picture of Greene's unnamed hero we will have to test his words and his deeds for genuine cynicism or real defencelessness; and we do that by comparing and contrasting them with clearer examples of each. If the agent was really cynical would he have risked his life for a political ideal in the way he did? But wasn't he being ruthless throughout? But then again the death of his wife and his own imprisonment only temporarily hardened him as things turned out. And with each question we bring to bear other situations, possible personalities, and likely motives which direct our attention to some of the textual evidence in support of one description and some of the evidence in support of another. And getting a clear picture of what a fictional character is like may help us to understand actual people: we may be able to gain a deeper insight into the character of an actual person by bringing up Greene's confidential agent for comparison.

Of course, this sort of reasoning from comparisons is not restricted to the identification of personality traits or attitudes; for parallel questions arise concerning people's physical characteristics. R. W. Newell gives the example of the parents of a child disagreeing over

whether the child resembles his father more than his mother (28). And we might also think here of the example of the police artist who tries to construct a picture of the suspect by asking the witness whether he had this or that sort of nose, whether his eyes were set apart this far or were small and deep-set like these, and so on through the basic elements of the human face. We even ask ourselves whether Rover looks happy or bemused. Or consider any sort of case where what is required is perceptual assessment of one sort or another. Say the problem is that of arranging several smallish paintings on a wall without them looking cramped, or too spread-out. Or getting the right colour for the wallpaper to match, or at least not to clash with, the curtains. Or putting the right amount of salt into the stew. In each of these, rather mundane, examples the procedure we employ is that of comparing and contrasting, of trying our samples and of trying to get it right.

Now it might be objected that these last examples especially, and some of the earlier ones as well, are not really examples of problems and questions for which reasoning is possible. But it is only if one thinks of reasoning as something which requires a linear arrangement of premisses and a conclusion that one will fail to see how one can reason in cases like these. Put simply, one reasons here by giving reasons, reasons which support one rather than another way of solving the problem, answering the question or removing the perplexity. If I claim that the best arrangement for the paintings would be to put them all, frame-to-frame, in the bottom, left-hand corner of the wall, or all clustered in the middle, or one in each corner, reasons can be given against the desirability of each of these proposals. 'If you put them all together like that they will look like you wanted to cover up a

hole in the wall', 'clustered in the middle like that will give the impression that what you really wanted was one large painting, not several little ones', 'one in each corner would be like having a flower arrangement with a rose in the bedroom and a tulip in the kitchen'--these are all reasons, they count against the proposals. Giving them is reasoning toward the best way to arrange the paintings.

But with questions about the best way to arrange paintings we seem to have come a long way from questions about foreseeability or reasonable care. Yet, although the content of these questions differs--and although we obviously take some of these more seriously than others--still they are on all fours in another respect: In each of these cases there is a similar sort of puzzlement which gives rise to the question. We are troubled by our inability to put the facts which are open to view into a pattern which will enable us to see the situation in the proper light. I have been saying that the general perplexity here is a conceptual one; but that is just another way of saying that we are initially unsure whether the facts before us warrant a particular description. It is the legacy of Wittgenstein and Wisdom that philosophical problems and questions are of just this sort. I shall want to say something about this view later--after, that is, I have finished trying to give an example in favour of the view. For the present, I wish only to claim that the source of moral perplexity, for example, the basis for the incompatibility which exists in moral conflicts, are questions of the sort we have been considering in this section.

Examples abound. Suppose we restrict our attention for the moment to the single moral problem of abortion. Even here the number of questions which are relevant to the question of whether it is morally

permissible for a pregnant woman to have an abortion, or for a doctor to perform one, is amazingly extensive. Setting aside questions which have to do with the ancillary issue of the morality of having laws which govern the practice, we might survey the relevant questions in brief.

There are questions about: (A) The status of the fetus: 'Is a fetus a person?', 'At which stage in the gestation process is a fetus a person?', 'Is the fetus only a potential person?', 'Is it an innocent person?', 'Can we properly distinguish between a person and a (mere) human organism?', 'Is the fetus just a part of the mother's body?', 'Is it a parasite on the mother?', 'Does the fact that the (pre-viable) fetus is totally dependent on the mother change the fetus' status as a person?'. (B) The right to life: 'Does the fetus have a right to life?', 'What reasons do we have for thinking that anything has a right to life (e.g., capacity to envision a future for itself, the ability to experience its environment, self-consciousness, viability, the ability to recognise a right to life, and so on)?', 'Can a potential person (or a mere human organism) have rights of any sort?', 'Is the right to life of a fetus comparable with the right to life of the mother?'. (C) Ownership: 'Does the mother own the fetus, as a part of her body?', 'Can one always do whatever one wishes with a part of one's body?', 'Does the fact that the fetus is totally dependent on the mother, or is a parasite, make the fetus the mother's property?', 'Does being willingly impregnated make the mother responsible for maintaining the life of the fetus, whether she owns it or not?', 'Does the father share ownership with the mother?'. (D) Special circumstances: 'If continued existence of the fetus is a direct threat to the mother's life is abortion just a form of self-defence?', 'Should prevention of emotional trauma (cf. the legal notion of nervous shock)

count as a form of self-defence?', 'If a pregnancy is the result of rape, does the status of the fetus change?', 'Would it be cruel to allow a grossly deformed child, or an unwanted child, to be born?'.

All of these questions, and many more besides, are obviously relevant to the problem of understanding the concept of abortion, and so to the central questions 'Is abortion the killing of an innocent person?', or alternatively 'Is abortion the surgical removal of a morally insignificant part of the mother's body?'.

The moral problem of abortion is not a unique sort of moral problem in this respect. It is a representative and indeed paradigmatic moral problem, the sort of problem which typically and regularly leads to moral conflict. Moreover, most perplexing moral problems, like the problem of abortion, have several facets to them. They rest on a large number of relevant sub-issues some of which may present themselves as separate moral difficulties. The questions I have listed above are all interconnected in various ways. Answers to some of them will count as reasons which are relevant to the answers to others. But there is a parallel here with the question of negligence in the court of law. There is with the question of negligence sub-issues which set out different lines of approach to the main question. We must consider the intentions of the agent, the nature of the harm done, the foreseeability of harm, carelessness, the presence of special responsibilities, and so on. The conclusions we reach concerning each of these separate and interconnected issues count as reasons in favour or against the judgement that the case was one of negligence.

So too with the question of abortion. The questions I have listed are representative of the questions which have been asked in the on-going

debate on the moral permissibility of abortion. We have here questions about the concepts of person, potential person, human organism, part of the body, viability, ownership, dependence, parasite, self-defence, cruelty, not to mention the concept of abortion itself. Indeed, with abortion the situation immediately becomes more complicated still as the single question of the moral status of the act begins on examination to fragment into several: 'Is abortion in the case of a clear threat to the mother's life permissible?', 'In the case of rape?', 'When either the mother or the child, but not both, can be saved?', and others which pick out special circumstances of the act.

It is clear enough that the standard treatment of moral reasoning fails to isolate the proper locus of the reasoning for complex problems like abortion. The real work is required at the level of the questions we have listed, not at the level of a question like, Are abortions morally permissible? And as with legal questions, deductive reasoning is relatively unimportant here. If we knew that all abortions were murders, and that murder is not morally permissible, then we would have little difficulty in reasoning to the conclusion that abortions are morally wrong. But it would not be an exaggeration to say that this particular inference, although logically impeccable, has nothing whatsoever to do with moral reasoning. For who would ever bother with such an argument? Imagine someone being perplexed by the moral status of abortion who was already convinced that abortion was a kind of murder. If abortion could be seen as murder, if the weight of the evidence pointed to the correctness of that description, most (if not all) of the reasoning and reflection the problem requires would have already taken place. Characterising the paradigmatic moral question requiring reasoning as a question about what

is good, bad, obligatory, permissible or forbidden, is like saying that the paradigmatic scientific question requiring reasoning is whether the Special Theory of Relativity is true, or that the paradigmatic mathematical question requiring reasoning is whether Goldbach's Conjecture is true (cf. A. Gewirth's "Positive 'Ethics' and Normative 'Science'").

The object of reasoning about the status of the human fetus, its innocence, its and the mother's rights is to see what sort of case can be made with regard to the general question of the morality of abortion. The general question is only made concrete, and the moral problem involved is only made intelligible to us, if it is reflected in particular questions about the morality of particular cases of abortion. No one would presume that he had said anything informative if he simply asserted that abortion was (or was not) morally permissible. The general claim makes for an ear-catching preamble to the informative moral pronouncement of what it is about the practice of abortion which makes it morally permissible or not. But then we are talking about features and circumstances of particular, real or hypothetical, cases. Indeed our questions tend to place the emphasis on features and circumstances which can then be grouped into various scenarios which set up different instances of abortion. If we agree that the fetus has the status of a person, or a potential person, we may then agree to set that parameter constant and vary some of the other features: if the fetus, as a person or potential person, has a fundamental right to life is that right such as to always outweigh the mother's right to her own body? And if so, what if the life of the mother is threatened, or what if we are forced to choose between saving the life of one or the other? And the argument continues. But in this dialogue the concept of abortion is developed and given life by the

various relevant features of the act and surrounding circumstances which are brought out by the questions we have suggested. This involves bringing up case after case, shifting the features, bringing now one now another feature into prominence, contrasting and comparing.

The object of this way of reasoning is to become clear about what, morally speaking, the act of abortion is. Something comparable happens in the common law. When the court deliberates over a particular example of putatively negligent behaviour the point is to see that behaviour against the background of the concept of negligence, in the light of past cases of negligence or (as often happens in judicial decisions) in the light of possible or hypothetical cases of negligence. The judge operates with, and presupposes, a common background of beliefs about what negligent behaviour is; he does not have to start at the very beginning reasoning from utterly blatant cases. The judge assumes the sort of knowledge of the concept which is common knowledge--the basic familiarity with what it means to behave negligently which must have already been employed to bring the case to trial. But to become clear about the instant case is just to see it as a case which is sufficiently similar, or sufficiently dissimilar, to an interconnected body of cases of negligence. But then, to see particular cases of abortion as cases of murder, cruelty, the exercising of a legitimate right or the surgical removal of unnecessary tissue is to see the concept of abortion in these ways. Particular cases of abortion fade into the general concept of abortion, while the general concept dissolves into the particular cases.

The heart of the reasoning in the case of a moral problem like abortion involves essentially descriptive concepts and questions which shift our point of view from relevant feature to relevant feature. But

it is natural at this point to object that, if it is true that the real work is done at the level of questions which involve descriptive or factual concepts, and if it is true that for such conceptual questions deduction is a relatively unhelpful rational tool, then why should not inductive reasoning be the rational tool we need. Inductive reasoning, after all, is the sort of reasoning which aims to expand our knowledge of the facts. It does not bring forward necessary inferential connections which are inappropriate when the problem concerns the exploration of the boundaries of concepts. And, as Mill thought, it is reasoning from the particular to the particular. We have agreed, it should be recalled, not to follow Strawson and extend the boundaries of induction so as to coincide with all forms of non-deductive reasoning. But isn't what I have been describing as reflection just induction in the strict sense? To meet this objection, and to prepare the way for a discussion in depth of the sort of reasoning which, so far, I have only been hinting at, I should spend some time supporting the claim that the sort of reasoning which is typified in legal, moral and indeed in most everyday cases is neither deductive nor inductive.

Footnotes

Chapter III: Reasoning from Particulars

(Page 79)

¹The term 'working logic' is S. Toulmin's and is found in his The Uses of Argument. He distinguishes working logic from what he calls 'idealised logic', or what I have been calling formal deduction and induction, i.e., logic. The term has for Toulmin several associations which appeal to me: working logic is rational argumentation, dialectics, the reasoning which in a judicial context yields reasoned decisions--in short, exploratory reasoning. Toulmin, however, is engaged in a project with which I am not in complete agreement. He is primarily interested in showing that idealised logic, the analytical syllogisms, is not only not the paradigm of reasoning (a point I am in total agreement with), it is also not a model of reasoning at all:

Rational demonstration is not a suitable subject for a timeless, axiomatic science; and, if that is what we try to make of logic, we are in danger of ending up with a theory whose connection with argument-criticism is as slight as that between the medieval theory of rational fractions and the 'music' from which it took its name. (147.)

Moreover, Toulmin's paradigm of reasoning consists of a set of informal argumentative techniques which are, if not straightforwardly deductive or inductive, then procedurally deductive: Toulmin is interested in a kind of procedure which has been called 'premiss guessing', setting up the skeleton of a deductive argument and trying to determine what premisses would be required to be able to deduce the desired conclusion. Still, Toulmin lays out an approach to the nature of reasoning with which I am very sympathetic. In particular, his suggestion that 'substantive' reasoning is paradigmatic and field-dependant is closely aligned with the suggestion I wish to make below that reasoning is topos-orientated.

(Page 86)

²This suggests a third kind of justification, justification by failure to falsify. Karl Popper has argued vigorously over the years that the basic methodological technique in science is not inductive in nature (a form of reasoning which Popper denies has legitimacy). Rather, one proceeds by attempting to falsify a scientific hypothesis by devising crucial experiments and the like. Each unsuccessful attempt to falsify a hypothesis counts, Popper argues, as confirmation (or in his terminology, corroboration). Extending this suggestion, we might think of unsuccessful attempts to falsify the principle of non-contradiction as counting as reasons for thinking that this principle is secure as a canon of rationality.

But this process is on all fours with the process of justification which I am suggesting is primitive. Attempting to show by counter-example

that a principle or general assertion is false or less likely is an attempt to test the principle by considering the justifiability of instances of it. Justification by a consideration of counter-examples is justification by a consideration of cases. The claim that something holds in general is always a promissory note, and that note is partially paid when instances are brought forward to show that that which has been claimed to hold in general holds in particular cases. Moreover, a principle is also a promissory note to the effect that particular counter-examples will not be found. In the end, as I argue in the text, the source of justification is particular examples, examples which count as reasons for thinking that something is so, and examples which count as reasons for the opposite judgement.

(Page 116)

³Of course, to call the sort of question at issue here a conceptual question is not to deny that it is a legal, or jurisprudential question as well. One might say that the question is a legal-conceptual question, where one would mean by that both that (i) the question of whether, in this case, the defendant should have foreseen harm is a question about part of what it means to exercise reasonable care; and (ii) the issue of reasonable care is to be asked and answered in a legal context, from a legal point of view. Questions about whether people have exercised reasonable care in their actions are, from the legal point of view, questions of fact, questions over which juries deliberate. But given the context in which the question is asked, the answer, the judgement arrived at, reflects a legally relevant understanding of what reasonable care comes to, an understanding which it is the job of the judge to provide in his instructions to the jury. In short, though obviously a legally significant question, the question of foreseeability (and others that follow in the text) is still a conceptual question.

(Page 125)

⁴Lord Denning, speaking of the degree of cogency which the evidence on a criminal charge must reach before the accused can be convicted, has declared:

That degree is well settled. It need not reach certainty, but it must carry a high degree of probability. Proof beyond a reasonable doubt does not mean proof beyond the shadow of a doubt. The law would fail to protect the community if it admitted fanciful possibility in his favour, which can be dismissed with the sentence 'of course it is possible but not in the least probable' the case is proved beyond reasonable doubt, but nothing short of that will suffice. (Miller v. Minister of Pensions (1947) 2 All E.R. 372, at 373-374.)

(Page 128)

⁵Graham Greene, The Confidential Agent (Penguin, Middlesex, 1971), 64.

CHAPTER IV

EXPECTATION REASONING

There are many ways of distinguishing deduction from induction, some better than others. One way, better than some, is this. With deduction everything is out in the open: we have a clear idea of what the inferential 'move' looks like, where it starts from and where it ends. But with induction, although we are usually well acquainted with the premisses--or, at least don't have many epistemological worries about them--the conclusions are about things and events which are, for one reason or another, out of reach. We take chances with induction; we can't go wrong with deduction. The problem is, therefore, how to characterise induction in a way which retains its distinctive character without sacrificing its rationality, in a way which shows why it is a form of reasoning without making it seem more like deduction than it is. The way I have chosen to approach the problem here rests on the claim that induction is expectation reasoning, a way (and perhaps the way) of reasoning to expectations. The claim that induction is expectation reasoning is a thread which runs through this chapter and connects the two tasks which I have set out for myself: (a) to get some idea of the kind of reasoning inductive reasoning is, and how it satisfies the two requirements of reasoning: security and epistemic advance; and (b) to show that reflective reasoning is not a species of inductive reasoning.

Induction--Utility and Legitimacy

There is with inductive logic, as there is with deductive logic, a distinction to be made between formal and informal modes of reasoning. Moreover, I shall soon be arguing that as with formal deductive reasoning, formal inductive reasoning acquires whatever legitimacy it may have from our recognition of whatever legitimacy informal inductive arguments have. It is necessary to speak here of 'whatever legitimacy' informal inductive arguments have so as to leave open at the start the question of whether it is possible to vindicate induction, that is, the problem of the justification of induction. To be sure, my interest here is not to solve this problem; the best I could hope for is to shine a bit of light on it. But, whatever light I do shine on this venerable philosophical problem will in the end be the light generated by two suggestions. The first is just that the problem of justifying induction is the problem of justifying our ordinary and informal inductive arguments. And the second is that the problem of justifying induction is just the problem of convincing ourselves that inductive reasoning satisfies the by now familiar legitimacy or security requirement of reasoning. I shall have more to say about vindicating induction later; but for the present I want to take a look at informal inductive reasoning before passing on to a brief discussion of the technical complexities of formal inductive logic.

There is little doubt that inductive logicians at times make exaggerated claims about their pet form of reasoning. The comment is frequently made that inductive reasoning is not the exclusive property of the scientist or the statistician since it provides the basis for almost all of our most familiar expectations, our beliefs and our actions. G. H. von Wright, for example writes: "When I assume that the same food

that nourished me yesterday will nourish me today, or that if I put my hand into the fire it will hurt, I am making inductive inference" (1). And S. Blackburn has recently cited as the sort of belief that is the result of inductive reasoning "the belief that my typewriter will not suddenly become weightless, change colour, or talk to me" (2). Now the claim that inductive reasoning is ubiquitous in this way does seem to follow from the standard picture of inductive reasoning as reasoning we employ when we suppose that the future will be more or less like the past or that the yet to be observed will be substantially like the already observed. But there is something wrong here. Surely we must distinguish between beliefs which are inductive conclusions and beliefs which are assumptions or expectations of uniformity and continuity. Everyday life would be a very complex affair indeed if every belief about the on-going existence of tables, chairs and typewriters and every belief about every regularity was the conclusion of an inductive argument. This is surely just an example of the philosopher's special disease of over-rationalising everyday affairs.

Suppose we call the kinds of beliefs which von Wright and Blackburn cite inductive beliefs and refuse in general to call them inductive conclusions. Doing so we would of course acknowledge that any inductive belief could have been an inductive conclusion, or indeed that if we found it necessary to do so we could conjure up a relatively plausible inductive argument for any justifiable inductive belief. But we can acknowledge these things while consistently maintaining that an inductive belief is not ipso facto an inductive conclusion. When I reach in the dark to turn off my alarm clock I assume that it will be where it has been in the past; but it would be very odd to say that before I reach,

before I even know why I am reaching, I must engage in some sort of reasoning. Sometimes we just sit in chairs, place books on tables and use typewriters in the normal way without giving the matter a second thought, without, that is, even assuming or expecting anything about the continuing (normal) existence of these things (cf. Wittgenstein (4), e.g. § 7).

Now Hume, on his way to the position of general scepticism about inductive reasoning, made a similar point:

When a child has felt the sensation of pain from touching the flame of a candle, he will be careful not to put his hand near any candle; but will expect a similar effect from a cause which is similar in its sensible qualities and appearance. If you assert, therefore, that the understanding of the child is led into this conclusion by any process of argument or ratiocination, I may justly require you to produce that argument....(Sect. II, Part II)

And Hume thought that no such argument could be forthcoming, given the limitations of the putative reasoner. Specifically, Hume claimed that the child--or for that matter a "brute beast"--does not reason from a previously learned general principle concerning the uniformity of nature. His experience with candles is necessarily prior to his understanding of the general principle. Hume, of course, went on to question the justifiability of all inductive beliefs and conclusions, but on this particular point he was surely correct: not every inductive belief is an inductive conclusion. Indeed, most of our inductive beliefs are unreasoned beliefs.

But they are not necessarily unreasonable beliefs, beliefs for which no inductive argument could be given. We could reason to the conclusion that the same food which nourished us yesterday will do so today, or that fire will hurt, or even that this typewriter will not suddenly become weightless--and there may be special contexts in which such arguments have a point or are even required. But these contexts are

special. We do not reason inductively to the vast majority of our inductive beliefs; we just have them. But then, inductive reasoning is not any less ubiquitous, or any more the exclusive property of the scientist or the statistician if it is the sort of reasoning we employ to arrive at that subset of our inductive beliefs which are genuine inductive conclusions.

What I am getting at is just this. We should look at inductive reasoning not as some sort of foundation for all of our inductive beliefs, but as a rational process which we deliberately employ to come to inductive conclusions. Induction has to do with our expectations; it is expectation reasoning. But it is not reasoning which makes an expectation an expectation. If we look at inductive reasoning in this way--as a deliberate process of reasoning--we shall be in a far better position to assess induction as a form of reasoning.

So, what is this form of reasoning? The simplified textbook answer is that induction is reasoning about expectations in this sense: the conclusion of an inductive inference is an expression of an expectation that uniformities observed in the past and the present will be representative of unobserved uniformities. Now although this is a bit misleading, it will do for a start. The textbook answer has the advantage of bringing one obvious and central feature of induction out into the open. Inductive reasoning enables us to extend the boundaries of our knowledge by allowing us to infer a hypothesis about what is presently, for one reason or another, inaccessible to us. And this feature of induction surely suffices to secure the utility requirement. Indeed, induction is particularly well suited to the requirement of epistemic advance since the conclusions which are inferred must go beyond what is

stated in the premisses of the argument. Of course, with the inconclusiveness of inductive conclusions comes a vulnerability to doubts about the justifiability of the inferences. One must pay a price for the chances one takes. It would seem that conclusiveness and epistemic advance are in tension. From the start, induction and deduction present opposing problems as far as our two requirements of reasoning are concerned. While deduction appears to stand secure on the legitimacy requirement while faltering on the epistemic advance requirement, induction appears to cinch the epistemic advance requirement while opening itself up to doubts which are directed toward the legitimacy requirement.

Still we can talk with a good deal of precision about the form which induction inferences can take. Since induction is expectation reasoning, it is possible to roughly separate inductive inferences into two groups on the basis of the logical character of the expectation which one concludes with. Inductively, we either reason to generalisations or predictions (and the link between generalisations and predictions is deductive). The cash value of an inductive inference is obviously the singular events which are predicted, since generalisations are, in principle, unverifiable. It is standard, and convenient, to divide inductive inferences into three very basic varieties:

(A) From the fact that something a was P at a certain time t_1 (or for some temporal interval), we infer that a is P for all times (generalising) or that a will be P at some future time or some future temporal interval (predicting).

(B) From the fact that a finite number of things of kind P are also of kind Q, we infer that all things of kind P are Q (generalising), or that an unobserved thing of kind P will be Q (predicting).

(C) From the fact that the ratio of the number of observed P's which were also Q to the number of observed P's was m/n , we infer that this ratio holds for all cases of P (generalising), or that the probability of a future P being Q is (approximately) m/n (predicting).

I want to set aside for the moment inferences of the third sort, (C), inferences which make use of the quantitative notions of frequency of occurrence and probability. Of the remaining inferences, two deserve separate mention. These are the two generalising inferences of varieties (A) and (B), inferences which traditionally go by the name of 'enumerative induction'. Now induction by enumeration was thought by the classical writers on induction to be the paradigm, if not the only correct form, of inductive reasoning. If the stronger of these two claims were true then all (correct) inductive inferences would be of the form of a finite list of premisses, 'This P is Q', 'That P is Q',... (or: 'a was P at t_1 ', 'a was P at t_{i+1} ',...) and a general conclusion, 'All P's are Q' (or: 'a is (always) P'). These days, however, inductive logicians agree that it is neither a necessary nor a sufficient condition of the propriety of an inductive inference that it be enumerative in form. It is not a necessary condition since, within a sufficiently determinate branch of science like physics or chemistry, it is often correct to inductively infer to a generalisation or a law from a single instance (see, for example, L. J. Cohen's extended discussion of this point in his The Implications of Induction, 152 ff.). And the thesis that the enumerative pattern is sufficient for inductive correctness has lead to certain counter-intuitive results, the most prominent among them being Nelson Goodman's notorious 'grue-bleen' problem. Moreover, given the prevalence of predicting inductive inferences in scientific and everyday contexts, it

would be equally incorrect to say that the enumerative pattern is paradigmatic of inductive inferences.

Nonetheless, it has been argued that ultimately the support which a predicting inference may have depends on there being a prior belief in the justifiability of a generalising inference from which the prediction can be deductively inferred. Von Wright has argued that predicting inductive inferences are "logically secondary" to generalising inferences (11 ff.). But it is not at all clear that this is true. For even if we set aside scientific examples, it is just not the case that when I predict that something or other will happen in the future I always deduce my prediction from a prior generalising inductive inference. If I argue that since it is now 12:05 and I haven't left the house I will miss the 12:15 Oxford bus (on the grounds that I have missed that bus twice before by leaving the house after 12:00), I am neither interested in nor relying upon the general proposition (which is very likely false anyway) that whenever I or anyone leaves the house after 12:00 that bus will be missed. More often than not in everyday arguments the generalisation which would be induced is either unlikely or inappropriate whereas a particular prediction based on the same evidence is both justifiable and appropriate. If it has been my experience that typewriter repairmen in large cities are unreliable, I may feel that it is good policy this time around to have my typewriter repaired in a village without wanting to claim that every typewriter repairman in every large city is unreliable. Indeed, I may have very good reasons for knowing that that general claim is not only rash, but false while consistently maintaining that it would be a good thing this time for me to take my typewriter somewhere else.

But the suggestion that predicting inductive inferences are logically

secondary is objectionable for another, and by now familiar, reason:

The inductive support we have for a generalisation is neither more than, nor different in kind from, the inductive support we have for particular instances of that generalisation. And this applies equally to those generalisations in science which are nomological in character. For even if the claim is that there is a law-like connection between two properties, the confidence we have in the general, and necessary, claim that whenever property P is present property Q is also present is a confidence which is ultimately dependant upon our recognition, in particular instances, of the necessary connection between the two properties. Since nomological generalisations are usually embedded in highly determinate systems, this dependence upon particular cases is not immediately obvious, for reasons which I tried to bring out in Chapter II. But, as I also tried to show there, ultimately the appeal is to particular cases, whether our concern is to show the over-all legitimacy of a logical calculus, the validity of a rule of inference or the internal consistency and explanatory power of a highly determinate, scientific system.

One last general comment about inductive arguments should be made before we move on to formal inductive logic. It is standard practice to include in the description of inductive inferences two parameters which are absent in the description of deductive inferences, namely the parameters of time and observation (direct acquaintance). It is tempting to mark off the difference between deduction and induction by making a great deal of this particular difference: deduction is 'atemporal', induction is rooted in the temporal continuum; deduction is concerned exclusively with relations between concepts, induction must always have both feet on the ground, immersed in observational data; and so on. Now I have followed

this practice of employing temporal and observational parameters in my descriptions of inductive inferences. And for a good reason: typically and paradigmatically inductive inferences are inferences from past and present observation reports to predictions about future observation reports. Still, it is important not to place undue emphasis on either the temporal or the observational elements of induction. These two parameters are obvious and paradigmatic examples of the epistemological dimensions along which inductive inferences move. But they are not the only examples of such dimensions.

Suppose, for example, we were to describe inductive inferences in the manner of (A), (B) and (C) above but instead of talking of past, present and future observations we talked about propositions about the verified or unverified, or propositions about what is accessible or what is inaccessible, or again propositions about what is acknowledged or what is unacknowledged. Of course, as it happens the temporal dimension offers the clearest picture of the sort of epistemic extension which characterises induction, and it would be very different, not to say silly and futile, to try to rip verified or unverified propositions out of the temporal continuum. We have little difficulty in understanding an inductive inference as one whose premisses are about what takes place at one point or interval in time and whose conclusion is about what takes place at another point or interval in time. That picture is a clear one, and there is no reason to displace it from its proper position, centre-stage. The observational dimension also presents an easily understood picture of the inductive move, although it fits more clearly into an empiricist epistemology where what is presently known is known by observation (and memory) and what is inductively inferred is what could have

been, or will be, observed. So both of these parameters fit the concept of induction nicely. Indeed, transcendental arguments there undoubtedly are which would show that the temporal and the observational epistemological dimensions are necessary presuppositions of the concept of induction.

But even if we grant the epistemological priority of the temporal and the observational dimensions, there is still room for the suggestion that there are other relevant dimensions which are usable, and are used. There are quite clear examples of inductive inferences which can either be described without any explicit reference to a temporal ordering of observations, or can be described in such a way that the observability of the states of affairs cited by the premisses and the conclusion is not a relevant concern. Examples of the first sort are easy enough to come by. Consider any recursive argument. With recursion we have an argument which is inductive if any argument is inductive but which requires only that the objects over which the induction proceeds be arranged in some sort of enumerable series, not necessarily temporal. We argue like this: Suppose the series $A_1, A_2, \dots, A_n, \dots$ can be enumerated in some manner (say, in mathematics by means of the 'less than' relation, or in logic by means of the subset relation). Then, if the first member of the series, A_1 , has a property \emptyset , and if when the n th member of the series has \emptyset the $n+1$ st member has it as well, then every member of the series has the property \emptyset .

It is possible, of course, to employ a recursive argument of this sort over a series which is temporally ordered and for which the observational dimensions is applicable (if you tip over the first domino it falls over and it hits the next one, and tipping over any domino in the line results in the one in front of it tipping over, then tipping the

first one over will make them all fall). Now it may be that proof by mathematical induction is a method of reasoning which rests on inductive arguments where the temporal and observational parameters are present. Nonetheless, there are difficulties with such a claim. For example, is temporal ordering epistemologically prior to, say, the 'less than' ordering of natural numbers? I am not at all sure how that argument would go; indeed it seems likely that something could be said for the priority being just the reverse. But however this question is decided, it is not directly relevant to the present point. Mathematical induction and recursive arguments generally are clear forms of inductive reasoning which operate successfully and intelligibly on dimensions other than those of temporal ordering and observation.

The other sort of inductive argument I am suggesting exists is one for which the whole question of whether the premisses or the conclusion are reports of actual or expected observations is not directly relevant to the logical status of the inference. These inferences are not at all mysterious, nor do they differ in form from other inductive inferences. These inferences are simply those for which the 'evidence' in the premisses is hypothetical or fictional, and intentionally so (I am not thinking here of inferences whose premisses are false, although we mistakenly believe them to be true). Suppose we ask ourselves this question: Would it be reasonable, from an inductive point of view, to conclude that a strain of wheat called 'X' was particularly suitable for growth in cold climates if there have been large yields of X ten seasons in a row in a region where the hottest it ever gets is 10° C? Could we answer this question if no yields of X have been recorded or observed, or for that matter if there is no such strain of wheat? Of course we could.

The question is about a particular case of inductive evidence--ten seasons in a row with large yields--and a particular inductive conclusion--X is suitable for cold climates. It is a question about a logical relationship between propositions. That is (recalling what was said earlier about deductive arguments with false premisses), our question here in effect asks 'Grant this much...now isn't it inductively reasonable that...?' Obviously we can and do reason this way, in an inductive fashion. But in this case the dimension along which the inference moves is not an observational dimension, it is a dimension which stretches from the acknowledged (or entertained) to the presently unacknowledged.

But how common, and more importantly how significant, are inductive inferences with premisses citing hypothetical 'evidence' and yielding hypothetical conclusions? It would not be an exaggeration to say that these inferences are more common than the textbook paradigms. Possible evidence, after all, forms a set which includes actual evidence as a proper subset. And it surely is the case that without such inferences the growth of the experimental sciences would be severely hampered. For scientists do not always proceed--as some philosophers of science would have them proceed--from empirical data on hand to theories and then to experiments designed to test these theories. Rather, scientists, like the rest of us, often argue from hypothetical premisses in order to see if they can determine what sort of empirical evidence would count as evidence in favour of a theory or general claim which, for other reasons, they think has a good chance of being supported by the facts. And scientists, like the rest of us, approach inductive as well as deductive arguments from various directions. Sometimes we begin with the conclusion and then go on to try to find the sorts of premisses which would inductively support, or deductively entail, that conclusion (a process

sometimes called 'premiss guessing'). But as far as the logic of the matter is concerned, premisses which cite hypothetical evidence are on all fours with premisses which cite genuine, observational evidence--the epistemological dimension along which the inductive inference moves may differ, but the logic of the inference itself remains the same.

It is important, then, to distinguish between the question of the logical status of a particular inductive inference and the question of its empirical value, or what is usually called the acceptability of the inference. Frequently (but by no means always), the conclusion of an inductive inference is an empirical claim, an expression of an expectation about what will be shown to be the case by direct observation. When the conclusion is offered as an empirical claim, the acceptability of the conclusion depends both on the logical status of the inference and on the truth, or initial probability, of the empirical claims contained in the premisses. Obviously, inductive conclusions based on hypothetical, fictional or counterfactual 'evidence' will not be very reliable as true empirical claims, just as deductive conclusions based on false premisses will not be very reliable as true claims. But the logical status of an inductive inference and the acceptability of the conclusion are two different measures of the value of the inference. They are different since the conclusion of an inductive inference may be unacceptable while still providing us with an epistemic advance. We may find out what would support the claim that X was suitable for cold climates (and construct experiments along those lines), but we wouldn't conclude that X was suitable because the experiments have not been conducted yet. The logical generation of probable empirical truths is one of the advantages of inductive inferences; it is, moreover, the primary advantage of such

inferences. But there are other advantages, advantages which have to do with the exploration of the possible. Exploring the possible is often the only means we have for coming to appreciate the actual when finally we encounter it. And reasoning inductively from hypothetical premisses is one way of carrying out this exploration.

The distinction between the logical relation between premisses and conclusions in inductive inferences and the acceptability of the conclusions of these inferences becomes much clearer as we shift our attention to formal inductive logic. For with formal inductive logic the basic informal inductive moves are reconstructed as quantitative probability relationships between propositions. But these relationships hold regardless of the source, the comprehensiveness or the factual truth or falsity of the evidence. In the case of formal inductive logic the acceptability of inductive conclusions becomes very much a separate problem, one which is not automatically solved by the quantitative machinery. Actually, it is important to see how little the introduction of the probability calculus changes the basic question of the status of induction as reasoning.

In most of its modern manifestations, formal inductive logic is a theory of when and how much different evidence inductively supports (or confirms) different inductive conclusions. Usually formal inductive logic is restricted to the sphere of scientific methodology and as a consequence the relation of inductive support (confirmation) is envisioned to hold between experimental evidence and scientific hypotheses. Inductive support is intended to be the analogue of the deductive relation of entailment: they both are seen as the fundamental inferential link which defines the over-all structure of the pattern of reasoning. The difference

between the two links is simply that whereas entailment either holds completely or not at all, inductive support holds to one degree or another. The support which evidence e gives to hypothesis h is complete if e entails h; but genuine inductive support is always something less than complete. e's truth does not guarantee the truth of h if e merely inductively supports h; rather, e's truth makes h more likely, more probable.

Thus, the relation of inductive support rather naturally gets explicated in terms of, and is frequently just identified with, the quantitative probability function $p(\underline{h}/\underline{e}) = r$ (read 'the probability of h, given e is r ', where r is a rational number $0 < r < 1$). The obvious attraction of explicating the inductive inferential link in terms of the notion of probability--or of identifying the two--is that unlike the intuitive relation of inductive support, the concept of probability comes complete with a well-established axiomatic calculus and a mathematical theory for calculating probability values. And this enables formal inductive logic to be modelled on statistics, an advantage given the vast amount of work that has already been done in that field.

Unfortunately, if we identify the inductive inferential link with the relative probability function $p(\underline{h}/\underline{e})$, or if we just say that inductive support satisfies the axioms of probability, we have merely substituted one problematic notion for another. For there are various ways of interpreting the notion of probability as well as competing theories giving different explanations of what the values manipulated so effortlessly by the calculus are supposed to measure. There is little doubt that the so-called statistical interpretation of probability (relative frequency in the long run) does match up closely with the inferences I have called (C) above. These are inferences from reports of the frequency of the

occurrence of a designated property in a finite sample population to a conclusion about the frequency of the occurrence of that same property for the whole population. Such inferences are clearly the staple diet of statisticians, the life blood of statistical analysis, and hence essential to most sciences. But they do not exhaust the range of inductive moves. Nor is this interpretation of probability suitable for the job of explicating the notion of inductive support since there are cases where e intuitively provides evidence which makes h more probable while not expressing statistically significant information (when, for example, e is the fact that a certain coin is biased in having two heads, while h is the prediction that the next toss will land heads).

In light of this, confirmation theorists invariably opt for the objective (epistemic or logical) interpretation of probability. This interpretation sees the probability conferred by e on h to be an a priori matter, a function of a logical relationship between e and h neither determined nor affected in any way by observations or any statement about how the things in the world in fact behave. (The difference between a statistical probability claim and an objective probability claim is often brought out in terms of the difference between asserting on evidence e that h is probable (statistical) and asserting that on evidence e h is probable (objective).)

Once the relation of inductive support is explicated in terms of (or identified with) objective probability, questions about the acceptability of particular inductive conclusions can quite easily be distinguished from questions about the logical status of the inferential link itself. So, in this respect, formal inductive logic incorporates a feature of inductive reasoning which I have suggested is part of our general understanding of

induction as a form of reasoning. This point also ties up with another made earlier about the necessity of distinguishing inductive conclusions from inductive beliefs. For if every inductive belief was an inductive conclusion it would be appropriate to ask, not only whether the belief is true, but also if there is an inferential link between premisses and that conclusion. But for the vast majority of our inductive beliefs (that the chair I'm presently sitting on won't carry me rapidly to the ceiling before I finish this sentence, for example) no premisses have been entertained: we just have these beliefs, we haven't reasoned to them.

The details of the various brands of formal inductive logic based on the notion of objective probability need not trouble us here. What we do need to consider, however, is how the introduction of probability affects the fundamental question of the justifiability of induction. What I want now to argue is that as far as this question is concerned, the introduction of probability leaves everything where it was before, and hence that, ultimately, the issue of the legitimacy of induction as a form of reasoning is a question which must be answered on the level of informal inductive arguments.

It is clear, first of all, that when we speak of the probability of a proposition what we are really interested in is the probability of the occurrence or non-occurrence of an event. If an inductive conclusion is a prediction about a particular case, then to speak of the probability of the prediction is just to speak of the probability that something will happen. And we are still talking about event-probability in the case of generalisations, although here the event in question is a change, i.e., the probability that all P's are Q's is just the probability that the event which would falsify the generalisation won't occur (hence:

$p((x)(\underline{P}x \supset \underline{Q}x)/\underline{e}) = 1 - p((\underline{P}x \ \& \ \neg \underline{Q}x)/\underline{e})$). This is most clearly seen in the ordinary, non-quantitative use of 'it is probable that', when we say that, for example, 'It is probable that if I leave the house now I'll make the Oxford bus'. The ordinary use of 'it is probable that' signals an expectation that something will take place. The quantitative version is an extension of this basic use and is most intelligible when there is a finite list of alternative outcomes of which we are antecedently aware, as in the probability that the next roll of this die will result in a six showing being 1/6. That the notion of probability is thus centrally connected to the expectation of something taking place accounts for the obvious relevance of talk of probabilities in connection with inductive (expectation) reasoning.

The probability calculus is a closed deductive system (optimally determinate) whose axioms are intended to reflect our basic intuitions concerning the informal notions of probability or likelihood. But the quantitative machinery of the calculus tends to obscure one very important fact about probabilities: Whatever evidence is relevant to a judgement about the probability of an event is also relevant to a judgement about whether that event will occur. The converse is also true since no evidence which is relevant to the judgement that an event will occur is irrelevant to the judgement of the probability of that event. This is just part of what it means to say that an event is probable or likely, and turning probability into a quantitative function does not alter the basic features of the concept.

Now this observation has lead S. Blackburn to claim that the whole notion of probability should be omitted from the discussion of the nature, and the justifiability, of inductive reasoning (97 ff.). Indeed, Blackburn

argues that, given the difficulties in interpreting the concept of probability in a manner which is suitable for an analysis of inductive support or inductive reason, the concept of probability constitutes a hindrance to the solution of the problem of justifying induction. It seems to me that Blackburn is right in this, but the argument he gives for the view can be put in a different, but I think more illuminating way:

The question which has to be asked of induction if it is to be correctly seen as a form of reasoning is, Given that it seems to provide epistemic advance, what reasons do we have for being confident that the conclusion of an inductive argument constitutes genuine epistemic advance? We should recognise this as the problem of showing that inductive reasoning satisfies the legitimacy requirement of reasoning. But, then, how does the analysis of the inferential link in the language of the calculus of probability help us to answer this question?

Suppose we say that an inductive inference is legitimate if and only if it is correct (or inductively valid) in the following sense: the probability of the inductive conclusion C (say, a prediction), relative to background knowledge B , is lower than the probability of C relative to B in conjunction with a set of premisses P , i.e., $p(C/P \ \& \ B) > p(C/B)$. But suppose we compare this with the following, non-probabilistic but minimal criterion of correctness:

- (1) An inductive inference is correct if and only if the set of premisses P , in conjunction with background knowledge B , gives us a better reason for thinking that the event which C predicts will occur than the reason (if any) which B alone provides.

Now the probabilistic criterion either says the same as (1) or something more than it (it could hardly say less, or something entirely different). If it says the same as (1) then the problem of justifying induction can be put either in the idiom of probability functions or in the idiom of 'better reasons for thinking'. And, given the inherent problems with the notion of probability, it might be best to stick to the less contentious idiom. If $p(C/P \ \& \ B) > p(C/B)$ says something more than (1), provides a clearer way of picking out correct from incorrect inferences, then what more could it say? Surely not that we know that the predicted event will occur. Inductive correctness certainly does not require foreknowledge. But then, the probabilistic criterion must be providing some stronger reason (short of knowledge) for thinking that C is an epistemic advance, stronger than thinking that the event predicted by C will occur. But we have already agreed that whatever counts in favour of an event being probable must also count in favour of thinking that it will occur. Therefore, the introduction of the notion of probability, despite the quantitative machinery which becomes available, leaves the question of justifying induction precisely where it was before the notion was introduced.

There are those who would strongly object to this by claiming that the quantitative machinery of probability theory surely adds something to the solution of the problem. But it is not clear what it adds, what it enables us to say which we couldn't say before. The fact that the standardly employed logical interpretation of the probability calculus results in a deductive calculus tends to mask the special problem of justifying induction by positing axiom-like 'inductive rules' which are used to deductively justify particular inferences. But the problem

remains since we obviously require reasons for adopting just those inductive rules among the infinitely many rules which could be adopted (vide W. C. Salmon, 50-52). But the whole history of the attempt to justify induction is marked by a related assumption--the need for a general rule or principle to justify induction as a form of reasoning. I would like to end this section by arguing that the question of the legitimacy of induction as a form of reasoning can only be begun to be answered once this assumption that a 'principle of induction' is required is itself challenged.

In The Problems of Philosophy, Bertrand Russell puts the assumption I wish to challenge in this way:

The general principles of science, such as the belief in the reign of law, and the belief that every event must have a cause are as completely dependent upon the inductive principle as are the beliefs of daily life. All such general principles are believed because mankind have found innumerable instances of their truth, and no instances of their falsehood. But this affords no evidence for their truth in the future, unless the inductive principle is assumed. (107)

The inductive principle Russel has in mind is essentially a rule of inference by means of which the various sorts of inductive inferences mentioned above would be validated deductively. Russell held that this principle must be assumed, as an axiom, and that no non-circular argument could possibly be offered for it. For the most part, attempts to provide a justification of induction have been attempts to argue for the principle of induction, to show why the general belief that the present will be like the past, or that the uniformities which we acknowledge are representative, is reasonable in itself. But why do we require a principle of induction in the first place? As far as our inductive beliefs are concerned, those which are not inductive conclusions, no such principle is needed. We might think here of Wittgenstein's remark in On Certainty,

The squirrel does not infer by induction that it is going to need stores next winter as well. And no more do we need a law of induction to justify our actions or our predictions.
(§287)

Since not all of our 'beliefs of daily life' are inductive conclusions, no justification for inferences which are not made is needed. But what of those cases where we do inductively reason. Russell himself notes that the general principles of science are believed because we have found innumerable instances of their truth, and no instances of their falsehood. And, as we have seen above, our confidence in general principles in the end is no more than the confidence we have in particular instances of these principles. So, once again, why do we need a principle of induction?

The reason why it has been almost universally believed that such a principle is required is obvious enough: most philosophers believe that justification by means of deductive entailment from principles is the only proper sort of justification there is. But consider this version of one part of Russell's proposed inductive principle: When P's have been found associated with Q's, and have never been found dissociated, the greater the number of cases of association the more probable it is that P's and Q's will be associated in fresh cases (cf. Ibid., 104). I suggest that the reason, the only reason, why we think that this rule is acceptable is that we have found that what is more probable tends to happen more frequently, and hence that we have found instances of the truth of the rule. So in this respect the principle of induction is on par with the scientific principles of which Russell speaks. There is an important difference between scientific principles and the inductive principle, however, and that is that the innumerable instances of the truth of the scientific principles are also instances of the truth of the

principle of induction. But it is important to note that Russell's rule does not provide us with any better reason for thinking that induction is a justifiable form of reasoning than the innumerable cases which exemplify that rule. But if the problem of justifying induction is viewed--as it is usually viewed--as the problem of providing a principle which would entitle us to have confidence in particular inductive inferences, and if that principle is thought to require more justification than that provided by our recognition of innumerable instances of its truth, then the problem will be insoluble since the proposed principle will always be vulnerable to legitimate counter-examples.

Inductive reasoning is, I have suggested, expectation reasoning. It is a form of reasoning which is an expression of our basic confidence in the uniformity and continuity of the world. That confidence becomes suspect as soon as it is articulated, because along with our faith in uniformity and continuity is our recognition that the unexpected happens from time to time. If our general confidence about the regularity of events is the belief that the exception is not the rule, it is also a belief that it is rational to proceed as if the world is uniform and predictable. For, ultimately, our confidence in uniformity and continuity is rooted in our pre-logical inductive beliefs, in the ways in which we regularly act and conduct our lives. And these inductive beliefs underwrite our confidence in the inductive conclusions which we arrive at.

So, if the problem of justifying induction is the problem of justifying, by means of an inductive principle, all of our inductive beliefs, then the request for justification is incoherent. It is reasonable enough to ask why we think a particular inductive belief is justifiable, for in particular cases we know how to proceed to answer

that request, we know what evidence is relevant and what not. Justification makes sense here. But if the request is for the justification of all of our inductive beliefs, all we can say is that this is the way we regularly act and conduct our lives. More justification could not be forthcoming. Justification ends here since the request for more serves to remove all grounds for being suspicious, for doubting the legitimacy of our beliefs. We have nothing to set this universal doubt up against in order to determine how we might proceed to remove it. (This is one of the lessons contained in Wittgenstein's On Certainty.)

But there is a genuine problem of justifying induction. The genuine problem is simply that of understanding what is meant by the phrase 'is a good inductive reason for'. It is the problem of devising grounds for the assessment of inductive inferences. The question here is not whether all of our inductive beliefs are justifiable, or whether we ever have a right to have confidence in an inductive conclusion. Rather the concern is one of clarifying the nature of the inferential link between premisses and inductive conclusion so as to better understand the conditions which allow us to have a right to expect certain things, when it is reasonable to be confident in an inductive conclusion and when it would be more reasonable to come to one inductive conclusion rather than another. It should be emphasised that for these questions the point made above about the usefulness of principles of induction is applicable. These principles are at most guidelines for the assessment of inductive inferences, statements which summarise in a compendious manner a vast range of inductive arguments which stand firm for us as examples of defensible pieces of reasoning. The problem of justifying induction, as seen in this light, raises many philosophical issues of

considerable complexity. But the problem is intelligible, and it is a problem which can only be successfully approached from an angle which provides the most familiar perspective on inductive arguments, that angle which illuminates the particular and informal inductive arguments which we are constantly employing.

Induction and Reflection

Some philosophers speak of induction with a hint of dissatisfaction in their voices. Inductive reasoning lacks the security that is part of the very nature of deductive reasoning; it lacks conclusiveness. There are no strict inductive rules that are exceptionless, that everyone would agree on or which could be used to silence the sceptic. Moreover, it would seem that what counts as good or sufficient for an inductive conclusion varies with the system and the subject matter. As a rule, statistical inferences within highly determinate fields like physics warrant more faith than sociological surveys which are modelled on the same, neutral, statistical formulas. Sometimes the result of a single, carefully conducted, experiment is viewed to be near-conclusive confirmation for a natural law. But when a political pollster calls a 85% positive response on a questionnaire a 'significant swing to the left', we have no qualms about discounting the conclusion on the grounds that the size of the population polled was too small or improperly randomised, or that the questions asked were ambiguous. But this is as it should be, and most philosophers have seen that inductive reasoning simply cannot be expected to live up to the level of security that one has with deductive reasoning. Of course, one cannot be too dissatisfied with induction since it is an essential tool for discovering, and confirming, new facts; it is a rational means for going beyond what we already know

or believe about the world. But is it the only rational means for doing this?

In Chapter III I have examples of reflective or deliberative reasoning which is often directed to questions about matters of fact and which proceeds from a consideration of particular cases and parallel examples. Reflective reasoning, I suggested, seems to be what is taking place in the legal forum, it is the 'working logic' of the judge. Moreover, the same sort of reflection is required in other areas, concerning other sorts of questions about the application of concepts. Moral questions, I claimed, required such a process of reasoning. But is all this talk of reflection and deliberation, and this suggestion that there is a form of reasoning which is case-oriented and non-deductive, just talk about induction? I have spent some time in the previous section characterising inductive reasoning and exploring the questions which were asked earlier about deductive reasoning, questions about legitimacy and epistemic advance. So we have a fair idea of the nature of inductive reasoning; enough of an idea to be able to compare reflective reasoning and inductive reasoning. And in this section, as a prelude to a separate examination of the nature of reflective reasoning, I want to bring out the differences between the two, and to show that reflective reasoning is not a species of induction, at least in the strict sense of the term.

It is important, first of all, to set aside one way of distinguishing reflective reasoning from both inductive and deductive reasoning which involves an issue which deserves separate and extensive treatment. This is the suggestion that non-deductive, non-inductive reasoning is in fact just persuasive rhetoric. There are two versions of this suggestion, one which condemns and one which condones and attempts to vindicate persuasive

rhetoric. The first is traceable to some comments Plato has Socrates making in the Gorgias and elsewhere about the vices of rhetoric and sophistry generally--the sort of pseudo-reasoning which was a marketable skill in Plato's Athens. The second is the relatively recent view that there is an essentially non-deductive, non-inductive form of reasoning which is demonstrably rational and justificative although the point of such reasoning is still to persuade and not to provide genuine support for the truth or falsity of the claims under dispute. I will be arguing in the next chapter that what Plato had in mind by 'rhetoric' can be distinguished from reflective reasoning (and, indeed, from what Aristotle had in mind by 'rhetoric') and when it is so distinguished it looks very much like Socratic dialectic, although there are still differences to be aware of. Similarly, of this recent advanced form of rational persuasion I will be arguing that it too can be distinguished from reflective reasoning in that the latter often succeeds in providing genuine epistemic advance while the former is thought never to be able to. But all of this will wait. Is there any other way of distinguishing reflective reasoning from inductive reasoning?

In his The Concept of Philosophy, R. W. Newell has offered an extended exposition of the views of Wittgenstein and Wisdom on the question of the nature, and central importance of reflective reasoning. In this discussion, Newell attempts to distinguish inductive reasoning from reflective reasoning in terms of what he takes to be the crucial difference between arguing from actual cases and arguing from actual or hypothetical (or fictional) cases indifferently. That is, in the case of reflective reasoning

the cases to which reference is made in support need not record anything which has actually happened, or is thought to have

actually happened, but only something which might have happened, even if it did not, and even if it is thought not to have happened. (65)

But, on the other hand,

an inductive argument cannot be based on possible, hypothetical or fictional cases, if its aim is to yield conclusions or hypotheses about what actually is so. (Ibid)

In the case of reflective reasoning, actual or fictional examples have "equal authority" as reasons in favour or against a conclusion; but induction must reject non-actual cases in favour of the records of actual observations (73).

There is little difficulty in seeing why Newell has this view. He is adopting, more or less uncritically, the Millian concept of induction. Newell quotes a clause from a passage in System of Logic which constitutes Mill's main thesis about the nature of induction. The complete passage runs as follows:

In every induction we proceed from truths which we know, to truths which we did not know; from facts certified by observation, to facts which we have not observed, and even to facts not capable of being now observed; future facts, for example; but which we do not hesitate to believe on the sole evidence of the induction itself. (II, 1, 3)

If inductive reasoning was (just) reasoning from facts certified by past or present observation to facts as yet uncertified, then Newell would be quite right in saying that an inductive argument cannot be based on non-actual, unobserved, cases. But this conception of induction is much too restrictive, and it is, moreover, just an offshoot of Mill's extreme empiricism. In the complete passage from the System quoted here it is easy to see that Mill is consciously equating 'truths which we know' with 'facts certified by observation'; and Newell carries on this characterisation of the premisses of inductive arguments by speaking of 'records of

actual observation'. It is unlikely that Newell would be willing to follow Mill's lead and suppose that every epistemic advance is a fact certified by observation, so why should he accept the Millian conception of induction which straightforwardly relies on that extremely dubious thesis?

The reason is, it seems to me, that Newell is overlooking a distinction I took pains to make in the last section, the distinction between questions about the logical status of inductive inferences and questions about their empirical value, or acceptability. Questions about the logical status of any sort of inference or argument are conceptual questions. They are questions about what it means for something to logically follow from something else, or for something to be a reason for something else.¹ But these questions can be answered without having to engage in an empirical investigation into the truth or falsity of the premisses or the conclusion of the inference or argument. Questions about the acceptability of inductive inferences rely both on the results of the conceptual inquiry and on empirical investigations and determinations of initial probability values (in the case of statistical inferences). But the issue of acceptability does not arise in those cases where inductive 'evidence' is intentionally hypothetical. There is little doubt that Newell has conflated the two kinds of questions that can be asked of inductive inferences. He writes:

In an inductive argument a person, in effect, bets on the existence of a striking analogy between observed cases and a case at issue on the grounds of the observed presence of less striking analogy. (73)

The bet Newell has in mind here is clearly the bet that future observations will substantiate an empirical similarity between a set of previously observed cases and an as yet unobserved case. As I have

already tried to show, this sort of bet is indeed what we usually make or have in mind when we reason inductively; but it is not part of the logical status of an inductive inference that the analogies we are betting on are always and only empirical similarities based on actual observations in the past and present and potential observations in the future.

One way to see this point more clearly is to ask what Newell would have to say about an inductive inference whose premisses were reports of fictional or hypothetical cases. Suppose we look at the sort of inductive inference which we have already seen to be essential to scientific theorising, an inference which serves as a guide to experimentation. I am suggesting that a scientist engaged in experimental work will conduct a 'thought experiment' by speculating upon the outcome of various experiments yet to be performed and reasoning inductively to the general conclusion or prediction he has in mind. In such a case, the inductive premisses are reports of results of experiments yet to be performed, reports of hypothetical experiments. But will Newell claim that the inductive inference the scientist uses to guide his experimental work is inductively incorrect, or not a genuine inductive inference? It would appear that he would be forced to say this, for on his view it is a matter of the logical form of inductive inferences that the premisses report actual observations. But if he were right on this, then an a priori issue--what sort of experimental evidence counts as inductive support for an inductive conclusion--is being argued to depend on a purely contingent fact--that the experiments proposed have not been carried out. And this is surely wrong. We certainly should be able to determine what sort of evidence will inductively support a proposition without having to consider the question whether we have the

evidence on hand or not. Moreover, the scientist who is arguing in this fashion is quite clearly aiming to infer hypotheses about what actually is so. He is inductively arguing, but not on the basis of actual observations.

What would be quite incorrect, if not bizarre, would be for a scientist to claim that inductive support has been established for his conclusion on the grounds that the speculated outcomes of hypothetical experiments constitute inductive grounds for that proposition. There is no such thing as 'armchair' experimental science. But what makes this claim the parody of scientific procedure it is has nothing to do with the logical status of the particular inductive inference involved. The mistake lies in treating hypothetical experimental evidence as actual evidence. Clearly, it is essential to the truth of empirical claims about the existence of a regularity that previous and past instances of that regularity be actual instances, and, ultimately, to have been observed instances. Thus, if it is the experimental method we are interested in, then there is an obvious and essential reliance upon observation: a scientific theory might be as logically tidy as anyone would wish but nevertheless false if predictions turn out not to correspond with the facts. In short, questions about the empirical value of inductive conclusions are questions which are centrally important to the scientific endeavour. The point I wish to make, however, is just that these questions are recognisably distinct from questions about the logical status of inductive inferences.

Now it might be thought that Newell has an obvious response to all of this. He could agree that the question of whether some inductive inference is logically correct can be answered without having to engage

in an empirical investigation into the truth or falsity of the premisses. Inductive inferences with hypothetical premisses may very well be correct inductive inferences. But, he might argue, he is talking about inductive arguments whose aim it is to directly yield conclusions about what, as a matter of empirical fact, is actually so, and for these arguments the empirical truth of the premisses is essential. Moreover, it is a matter of the logical status of such inferences that their premisses be empirically true. For ultimately the very rationality of the process of inductive reasoning rests on the observational. When in a particular case the premisses report hypothetical cases, we only recognise this as an inductive inference because the cases referred to could have been observed, they refer to empirical matters.

But here again there is an essential confusion between the logical status of an inference and the empirical value of an inference. For, as we saw in the previous section, not all inductive inferences yield conclusions about what, as a matter of empirical fact, is the case. I pointed out that recursive arguments are straightforwardly inductive although for most applications--namely those in mathematics and in logic--the conclusions reached have no empirical content, since if they are true at all they are necessarily true. There is little doubt that Newell conceives of inductive reasoning solely as reasoning about matters of fact (cf. 17). He properly objects to the claim that inductive reasoning is the only proper reasoning about matters of fact, but he fails to acknowledge or appreciate the fact that inductive reasoning has an important role, indeed an essential role, in mathematics and logic.² Nonetheless, it is true that the procedure of inductively reasoning to empirical conclusions depends upon the observational content of the

premisses, and hence that hypothetical cases may serve as premisses to inductive arguments yielding empirical conclusions only if the cases referred to could have been observed. But this dependence is not a matter of the logical status of inductive inference, it is a matter of the empirical value of inductive inferences. The logical status of an inference is not a function of the empirical content of the premisses or the conclusion.

Even if we restrict our attention to empirical inductions, Newell's attempt to distinguish between inductive and reflective reasoning fails. For, in the sense in which it is true that empirically-aimed inductive reasoning rests on the observational, it is also true that reflective reasoning rests on the observational. We need only look at the problem of reasoning about negligent behaviour. What I have suggested is that the details of the instant case of putatively negligent behaviour are set alongside descriptions of other sorts of behaviour which are analogous. But these descriptions, although not necessarily descriptions of actual cases, must nonetheless be descriptions of possible cases. In reflective reasoning about concepts like negligence, it is equally true that hypothetical or fictional cases may be used only if the cases refer to empirical matters, and hence are observable.

Newell's attempt to distinguish inductive from reflective reasoning, which is, in the end, an attempt to show that reflective reasoning is not a species of inductive reasoning, fails because it concentrates on what is an unimportant, and surface, difference between different examples of the two sorts of reasoning. If one picks the examples carefully enough, this difference looms large. Obviously the scientist who is seeking confirming evidence for a statistical generalisation is not going to spice up his experimental findings with the odd fictional experiment, and

obviously the literary critic who is trying to argue that a certain fictional character's behaviour is self-serving is not restricted in his choice of parallel cases to actual behaviour by actual people. But other examples can be chosen where the premisses of a perfectly correct inductive argument refer only to fictional cases or where the only parallel cases employed in a reflective argument were actual cases. But this does not mean that there is no difference between the two forms of reasoning.

What Newell has overlooked is the obvious difference in the orientation of inductive and reflective reasoning. The orientation of inductive reasoning is captured by the remark that induction is a form of expectation reasoning. But to get at the difference here it is important that we not put undue emphasis on the similarity in form between some inductive inferences and reflective reasoning or on the common use of hypothetical cases. Both of these features will, in the end, only show the important feature which both forms of reasoning share, namely that they are both non-demonstrative, non-conclusive forms of reasoning. Yet, once we take seriously the comment that induction is expectation reasoning, the differences which are there to be seen between cases of induction and cases of reflection will become apparent.

Inductive reasoning is paradigmatically orientated to the prediction. What is predicted need not be, at least in any straightforward sense, in the future or presently unobserved. If by means of a recursive argument I inductively infer the generalisation that every natural number greater than zero is the successor of another natural number, the cash value of my inductive conclusion is that any natural number you care to pick, as long as it is not zero, is the successor of another natural number. Or, if by means of a statistical inference I claim that 90% of those who have

a daily intake of not less than .003 micrograms of nicotinic acid will never contract pellagra, the cash value of my inductive conclusion is that for anyone you happen to pick, real or imagined, if he has the appropriate sort of diet, then the chances are 10:1 that he will never contract pellagra. Inductive reasoning, in a word, is motivated by the desire for new particular facts. It is directed to the expectation that something or other will occur, or will be shown to be the case, on the grounds that other things have occurred or are the case. The basis for the inference is, with respect to the conclusion reached, inconclusive grounds for certainty. We proceed from what might best be called clues to the expectation that something else will occur or is the case. That these clues do not provide conclusive support for the truth of the conclusion--but only, or so it is hoped, make the conclusion worthy of our expectation--is the primary difference between induction and deduction. That is, all deductive consequences are in some appropriate sense 'contained' wholly in their premisses, although not everyone can detect them there (cf. Wisdom (1), 197); yet, no inductive consequence is contained wholly in its premisses and no one could detect it there.

But if no one logically can detect an inductive conclusion in a set of premisses, it is hoped that one can detect in these premisses grounds for justifying the expectation that the inductive conclusion is true. But an expectation is the sort of belief which, logically, looks forward along one or another epistemological dimension to results of further investigation (and, ultimately, investigation of a direct, non-inductive sort). So much follows from what it means to expect something. If one expects that P, then one will proceed as if P is, or will be the case. Yet, one will be awaiting, anticipating or (if not these then)

taking for granted the results of further investigation. The degree of confidence one has in the results of investigations which have not yet taken place will depend on the sort of case one is considering. And the degree of confidence will determine whether we can talk of 'awaiting', 'anticipating' or 'taking for granted' future or further investigation. In the case of a mathematical induction, confidence is obviously high, so high that it would be a joke to suggest that we must await investigations 'far up' in the misty regions of the number line to see if each and every natural number is in fact the successor of another. But, of course, scientists do await and anticipate the results of further investigations--the slight deflection of light near the surface of Mercury, the decrease in the incidence of pellagra, or the results of an upcoming election.

And this feature of what it means to reason to an expectation is retained in the case of inductive inferences from hypothetical or fictional cases. Here one has to change grammatical mood, of course.³ If a scientist inductively argues from hypothetical growth rates to the conclusion that strain X would thrive in cold climates, the logical move is described subjunctively. Were this to be the sort of evidence on hand, then it would be reasonable to conclude that we should have confidence in further investigations confirming that X is suitable for cold climates. But this move to a subjunctive conditional is neither suspicious nor logically defective. It is a commonplace that natural laws entail (whereas mere empirical generalisations do not) subjunctive and counterfactual conditionals, making connections between hypothetical reports that a is a P and the prediction that it is also a Q.

We may say, moreover, that reasoning to an expectation is captured by inductive reasoning by means of a projection from premisses to

conclusion. This projection can be understood to consist of the application of one or more predicates to the individual(s) designated by the conclusion (the induced individuals, as we might call them). This projection thus expresses the expectation that these predicates will be applicable to the induced individuals. In the paradigmatic case of predicting inferences, the picture we can construct is something like this. We begin by isolating or presupposing a set of individuals--we do this by speaking about things of a certain specifiable kind. Out of this set we pick our sample. Claims about the applicability of a certain predicate to each individual in our sample constitute the premisses of the inference. But then, the applicability of this predicate is projected onto another individual from the original set, and the claim that this individual also has the specified property is our conclusion. Thus, for example, from the set of all (observed and unobserved) shrews we pick (or, more likely, we are presented with) a sample set of (observed) insectivorous shrews. We then project the property of being an insectivore onto another (unobserved) example of a shrew. An analogous picture can easily be constructed for the other varieties of inductive inferences.

But now other questions emerge: What is required for this projection of the applicability of predicates onto new, induced individuals to be, first, possible, and second, worthy of our confidence? Two things are required. The first is that the predicates involved must be projectible, in the sense of that term which N. Goodman in Fact, Fiction and Forecast has made popular. That is, the predicates involved must be extensionally definable, these extensions being built up from a core established from the sample set by an accretion of new 'discovered' individuals. (Being extensionally definable, or for short in what follows, being definable,

is the linguistic analogue of being, conceptually, determinate: I obviously assume a one-one relationship between concepts and predicates which express these concepts, such that a concept is determinate if and only if the predicate which expresses it is extensionally definable.) For the last example, this first requirement comes to this. The definition of the predicate 'is an insectivore' cannot change as it is projected onto future cases of shrews, although the extension of the predicate grows with the discovery of each new shrew. If this predicate did change, then we should not have confidence in our expectation that as yet unobserved shrews are insectivorous, in just the sense that the sample instances of shrews are. As Goodman has shown, the claim that there are projectible predicates is intimately related to the claim that induction is justifiable. The two claims are closely related because to deny each yields a similar sort of scepticism: to assert that a certain predicate is projectible is in effect to predict that the reasons we will have for applying it in future cases will be precisely the same reasons we presently have for applying it to present cases. In any event it is clear that for this requirement of projectability to be satisfied, the projected predicates must be at least assumed to express determinate concepts, i.e., there must be necessary and sufficient conditions for the applicability of such predicates to any case which should arise.

The second thing required for the inductive projection of predicates to constitute a procedure of expectation reasoning is this: it must be, in principle, possible to ascertain by direct and non-inductive means whether the projected predicate does indeed 'fit' the induced individuals. This implies that it must be possible to identify the induced individuals in an independent way--that is, without employing

the predicates which are being projected onto these individuals. Inductive procedures satisfy this requirement by means of an original specification of the set out of which the sample is drawn. Thus, the set of shrews is specified in advance, independently of the property of being an insectivore, and the induced individuals are identifiable as shrews before it is legitimate to ask whether the projected predicate 'is an insectivore' can be correctly applied to them. Moreover, for the inference to be truly inductive, and not deductive, being a shrew must not entail being an insectivore. But then, the predicates employed to specify the original set must be, or must for the purposes of that inference be assumed to be, definable predicates. We must be able to specify beforehand the necessary and sufficient conditions for something being an element of the original set. If this is not assumed, then the basis of the inference cannot support the projection made from it, and the prediction or generalisation which is induced is not a justifiable expectation.

Taking these two requirements together, it becomes apparent that inductive reasoning is built upon and presupposes certain decisions concerning the concepts which are employed in the inference. These are decisions about the determinate character of these concepts. Not only must we be able to say that our sample belongs to a specifiable set--and that new elements of that set can be identified by means of a property held in common--we must also be able to say that the sample unambiguously possesses the properties to be projected and that we can decide whether the induced individuals possess these properties (once these individuals become accessible to direct and non-inductive investigation). If the concepts employed in an inductive inference are not determinate and do not already fit into a determinate system, then, for the purposes of the

inference, they must be assumed to be determinate. As one reasons, one must refuse to acknowledge, at either end so to speak, the existence of borderline cases. The form of reasoning requires that premisses report the existence of an object which is of a certain kind and unambiguously possesses a certain property, and the conclusion predicts that an induced individual will, equally unambiguously, possess the appropriate properties. This means that on the linguistic plane inductive reasoning, like deductive reasoning, requires and presupposes the specification of necessary and sufficient conditions for the applicability of the predicates which are employed. Inductive reasoning is not reasoning about concepts, it is reasoning with concepts.

Now it might be objected that this last claim seriously oversimplifies, if not inductive reasoning itself, then the uses to which it is put in the overall scientific methodology. For the claim seems to be that inductive reasoning itself plays no role in the setting out of plausible criteria for the application of predicates, and hence no role in the construction of determinate concepts. But if we look at scientific methodology as a whole, the objection continues, we soon discover that inductive reasoning often yields changes in the criteria for the application of predicates. Von Wright, for example, has put this point in this way (40-53):

Consider the concept of phosphorus. Suppose we have some list of criteria C_1, C_2, \dots, C_n which at this (early) stage in chemistry we take to be singly necessary and jointly sufficient for something being phosphorus (say macroscopic properties such as colour, smell, taste and the like). By experiment we discover that pieces of phosphorus so identified melt at 44°C , i.e. that the melting point of phosphorus is

44° C? He answers that we have in fact two things we can say. We can either say that our inductive generalisation has been falsified, or we can say that the last examined substance cannot have been a peice of phosphorus at all. If we make this second claim--and as von Wright notes, scientists in practice often do proceed in this way--we are in effect denying that C_1, C_2, \dots, C_n are jointly sufficient for something being a piece of phosphorus. We then add another criterion--the melting point of 44° C--and the concept of phosphorus has thereby been changed. But it has been changed as the result of inductive reasoning.

It has indeed; but it is important to understand what sort of decision (for that is what is involved here) is required when a result such as this is found. If there is no doubt that the experiment itself has been properly conducted and if no reason can be found for doubting the result, then the decision is whether to treat the result as a falsifying instance of one inductive generalisation or as a counter-example to the sufficiency of a set of criteria which, up until this time, has been employed to define what phosphorus is. Now both the inductive generalisation and the sufficiency claim are connected in various ways to other beliefs we have in chemistry. The new result introduces a bulge in the fabric of our knowledge about phosphorus in particular and chemistry in general which must be smoothed out. But there are, at least logically speaking, countless ways of incorporating the new result into the fabric, although each of these ways will require some change to the whole system, spreading ripples in various directions. Some of these ways of incorporating the new result are ruled out straight away because they would require vast changes in the system, forcing us to amend highly confirmed physical laws or to reject other inductive

generalisations in which we have a great deal of confidence. The scientist must weigh one decision against another in terms of the overall effect to the system of beliefs which each decision would require if implemented. But, the decision here is taken against the background of a more general picture of what scientific methodology is like. Specifically, the decision how to interpret the new result about phosphorus is made on the basis of the assumption that there are necessary and sufficient conditions for something being a piece of phosphorus. What these criteria are is a matter for empirical investigation, including more inductive reasoning. But that the concept of phosphorus is determinate, whether we know what the criteria are or not, is not part of that decision; it is one of the things which make it intelligible to be making a decision in such a case. (Rather than deciding that there were no strict criteria for something being phosphorus, it would cause much less overall disturbance to the system to decide in a wholly ad hoc manner; to decide, for example, that the melting point of phosphorus changes every other Thursday.)

Thus, inductive evidence and inductive conclusions do affect the criteria by which certain concepts are constructed. But the change is orderly and governed by the ground rules of scientific methodology. It is important to notice that the reasoning which the scientist engages in to decide the question of how to treat, say, the phosphorus result, is not inductive reasoning, although, obviously enough, inductively derived results are employed. One cannot reason inductively to the conclusion that a particular borderline case is more appropriately called an X rather than a Y. Inductive conclusions may provide invaluable evidence one way or another ('If it really was an X then, chances are, it would also be a Z'), but inductive procedures are restricted by their very

nature as expectation reasoning from being the sort of reasoning one could engage in to reason about concepts. Induction--and deduction too for that matter--can bring to light inconsistencies and paradoxes which require adjudication and decision. But when the decision is about the boundaries of a particular concept, then induction (and deduction) are ineffectual: they are not the appropriate working logics.

In the last couple of pages I have been looking at inductive reasoning from the perspective of its role as expectation reasoning. I have been moving back and forth from conceptual points concerning what is required for a conclusion to be an expectation, and what is required for a conclusion to be worthy as an expectation, to points about the nature of inductive reasoning. I have been looking for a match. And there is little doubt that there is a match: inductive reasoning is clearly directed to the expectation that something will occur or will be shown to be the case; inductive conclusions are the sorts of claims about which it makes sense to say that conclusive confirmation is expected or anticipated, or would be expected or anticipated. Moreover, I have tried to show that the conceptual requirements for projecting predicates onto new individuals are just aspects of the notion of expecting these predicates to be applicable to the new individuals. Confidence in our expectations is conceptually tied to the projectibility of predicates. And this to a great extent depends upon the determinate nature of the concepts employed in the inference. Yet, one last general feature of inductive reasoning has still to be tied with the notion of reasoning to expectations.

This last feature is just the obvious one. It is the mark of an inductive inference that our confidence in the conclusion is increased

by positive instances and decreased by negative instances. This is the very heart of inductive reasoning--it is the inductive policy--and it is exemplified by every inductive rule. If more shrews are examined and found to be insectivorous, then we should be more confident than we were before that the next shrew we find will also be insectivorous. If m/n of observed A's are B, and p/n of observed A's are C, then if $m < p$ we should place our bets on the next A being a C rather than a B. If by polling 10,000 more registered voters the percentage of 'left-leaning' voters drops from 85% to 65%, we should be less confident than we were before that the left-leaning candidates will win the next election. And so on; the inductive policy provides the distinctive flavour of inductive reasoning. There are also two limiting cases where the emergence of a negative or a positive instance has immediate and obvious effects on our confidence in inductive conclusions, although not by virtue of the inductive policy. If we inductively generalise that all shrews are insectivorous, then the discovery of a single herbivorous shrew (actually herbivorous and really a shrew) reduces our confidence in the inductive generalisation to nil. And, on the other hand, if we predict that the very next shrew we see will be insectivorous, then if it turns out that that shrew does indeed have the appropriate eating habits then our confidence in the original prediction should be maximal. But these are limiting cases, since, after the new positive or negative instance is introduced as an additional premiss, the inference is no longer inductive: the fate of every inductive inference, when we have, so to speak, caught up with the projected image, is to be turned into a trivial deductive argument.

But what of this last feature of inductive reasoning, does it match up with what it means for an expectation to be justified? Is it the case,

that is, that the emergence of positive instances should logically affect the plausibility of an expectation, while the emergence of negative instances should logically affect the implausibility of an expectation? If we are talking about the two sorts of limiting cases, then there is a direct logical influence on plausibility or implausibility which results from the emergence of positive or negative instances. But, after the positive or negative instance is added to our stock of information, it is no longer appropriate to speak of an expectation in these limiting cases: what was an expectation is now fact or falsity. But if we turn our attention to other cases, the question becomes considerably murkier. What we have here is obviously the problem of induction all over again. The policy of increasing confidence with positive instances and decreasing it with negative ones is one way of understanding what it means to strengthen or weaken an expectation. But so is divination. Or a counter-inductive policy like: the more times a regularity is observed the less likely it will be exemplified again. But, however intuitively repugnant, and however unsuccessful, these non-inductive policies are, they are still policies that could be employed in determining the justification, or lack of justification of our expectations. In short, to baldly assert that the inductive policy is logically tied to the notion of expectation just prejudices the issue of the justifiability of induction. It seems to me that there are ways of indicating the logical connection between the inductive policy and expectation reasoning--ways I have hinted at in the previous section--but merely asserting that there is this connection isn't one of them.

To sum up: induction is expectation reasoning based upon a projection of predicates assumed to be definable onto new individuals

independently identified by predicates also assumed to be definable which adheres to a general policy of increasing confidence with positive instances and decreasing it with negative instances. Moreover, as expectation reasoning, it is reasoning toward a new fact direct and non-inductive confirmation of which is anticipated in the course of further investigation. Suppose we compare this picture of reasoning with the sort of reflective reasoning which a judge might engage in when considering an instant case. Three points of comparison can be explored.

(1) The point was made in the last chapter that the judge's reflection typically takes place after the facts are in and the stories have been told. The judge's task is to bring together the various ingredients of his decision--the facts of the case, precedents and statutes--and reason toward some sort of judgement about the instant case. The raw ingredients for his judgement, the source of his reasons, are all before him. But if this is the sort of setting in which reflective reasoning takes place, then the judge is clearly not reasoning toward an expectation about what will be shown to be the case once further investigation takes place. Judicial reflection is not expectation reasoning. The judge is not gathering information on the basis of which he can predict that the instant case will be shown, conclusively, to be a case of, say, negligence. If a judge's decision was a prediction, then the trial would not be concluded until whatever it was that was predicted to happen was observed to happen. But this is not how trials are conducted. It may happen, of course, that a judge will postpone his judgement until a sequence of events concludes, if the probable termination of those events will be relevant to the case. It is proper, in other words, for a judge to postpone his decision on a criminal trial until the victim of a violent crime either dies or shows positive signs

of recovery. But in such a case the judge is postponing his decision, not making a judgement as to what can be expected to happen.

This is not to deny that it is the nature of legal cases that crucial facts may be unknown at the time of the trial. In some cases crucial facts are just inaccessible, although judgements made in the absence of these facts are nonetheless rationally supportable. In the recent case of the Moorgate tube crash, for example, the driver of a London Tube train crashed his train killing himself and forty-six passengers. The question which became the centre of much controversy--whether the driver was negligent or was attempting to commit suicide--simply could not be answered one way or another for lack of information. Still, the judgement that the driver was responsible for the crash was made, and was rationally made on the basis of eye-witness reports that the driver was sitting upright and alert seconds before the crash and on technical reports that no mechanical faults were discovered. But this judgement is not a prediction that further investigation will show conclusively that the driver was responsible. The judgement came at the end of the investigation, it wasn't an anticipation of the results of further investigation. To be sure, if new and relevant information does surface then the case may have to be reopened. But reopening the case is not an indication that there was something wrong with the original reasoning, nor does it falsify, inductive fashion, the original judgement. New facts require a rethinking of the case, we must start from the beginning and treat the case anew.

(2) But then, since reflective reasoning in the law is not expectation reasoning, it doesn't have the same orientation as inductive reasoning. The motivation behind presupposing that the predicates

involved are definable is gone. Inductive reasoning requires, even as first approximations subject to revision, strict definitions of the terms which are employed, the predicates which identify the sample set and the predicates which are to be projected. But, as examples of reflective reasoning suggest, the point of reasoning about concepts is just to become clear about what does and what does not constitute an instance of a particular concept, when it is and when it is not correct to apply the associated predicates. There are fixed points of reference in reflective reasoning, the judge is not engaged in a process of reasoning which puts all of the relevant concepts up to the light at once. This is obvious. But the crucial, the key concepts involved are not assumed to be determinate from the start: the judge is not applying a formula for negligence, or for carelessness, or for that matter, for what constitutes a flood.

As I have tried to show deductive reasoning depends on principles of one sort or another; and inductive reasoning presupposes definable predicates. But if we either appeal to principles--in the deductive fashion--or if we presuppose definable predicates--in the induction fashion--then we are in effect assuming that questions about concepts have already been solved. Yet, in countless cases in the law, in morality, in philosophy and elsewhere it is just these sorts of questions which are crucial. The judge who considers the question of whether the instant case is one of negligence obviously has a fairly good idea of what a negligent act is like--he is familiar with precedents and other cases, he knows the guidelines laid down by the statutes, he knows what to look for, what to ignore and what to pay particular attention to. But none of the sources of the law presents him with a strict

definition of those terms which are pivotal. Nor does he come to the case assuming that there are definitions which future cases will show to be exactly right: he doesn't predict his own decision.⁴

A more sophisticated version of the claim that judicial reasoning is inductive in nature has been offered by L. J. Cohen. Cohen claims that assessments of support for the conclusions of legal arguments from precedent have the same logical form as assessments of experimental support for scientific hypotheses (vide, 155-171). He offers as the paradigm of an 'elementary legal hypothesis' the following: For any two people X and Y, if X bears some factual, legal or legal-cum-factual relation to Y, then X has a good cause of action against Y. The inductive evidence for this hypothesis, Cohen suggests, is judicial precedent, previous cases founded upon a precedent concerning the legal relationship in question. The judge reasons in terms of previous decisions in an inductive fashion by reading off the implicit prediction contained in the ratio decidendi of the precedent and judges accordingly.

Cohen admits that what he is offering is an idealisation, but he is not worried about that since he wants his account to stand as a recommendation and not a description. He is also sensitive to the objection that judicial decisions are judgements about what (legally) ought to be not about what will be. To this he responds that in the law inductive reasoning serves to provide, not predictions about how judges will decide future cases, but rather justification for these decisions, justification based on the principle that like cases ought to be judged alike. Cohen's induction, unlike O. W. Holmes's, is induction from one or more precedent to a reason in favour of a particular judgement. Sophisticated as it is, however, Cohen's account falls victim to a relatively obvious objection,

one which he considers but fails to appreciate.

The objection is that most legal reasoning is concerned, not so much with the truth or validity of legal rules, generalisations or 'hypotheses', but with the meanings of legally relevant concepts. In particular, legal reasoning is concerned with the meanings of key terms employed in the formulation of these generalisations, statutes and rationes alike. Thus, as Cohen puts it, "...this preoccupation with the determination of vagueness is a feature of legal reasoning which sharply differentiates it from any kind of deduction or induction" (165).

One might think that a response to this objection would have to show that the reasoning required to ascertain, in the light of new cases, the legal meanings of rationes or statutes is inductive reasoning. For, after all, the point of the objection is that deduction and induction are, with respect to the kinds of questions requiring reasoning, ineffectual. What Cohen supplies by way of a response is, unfortunately, not relevant to the objection. In effect, Cohen replies that if a ratio is reinterpreted, the inductive machinery can handle the change in the legal hypothesis. Schematically put, if the original ratio had the form 'All cases of F_1 and F_2 and F_3 are G ' (e.g., All cases of failure to stop completely at a stop sign while driving at night on crowded streets are cases of reckless driving), and a later judge decided that one of the antecedent conditions, F_2 (say, driving at night) is not material, then the original legal hypothesis can be modified in such a way that support for the first will also count as support for the modification. But, the interesting legal question here--the sort of question which Cohen merely ignores--is, Is the fact that someone was driving at night when he failed to respond to the stop light relevant to the judgement

that he was reckless? This question has to do with the meaning of the key term 'reckless', and Cohen has not shown how inductive reasoning helps to decide the question. In short, Cohen has pointed out how his formal machinery can deal with alterations in the ratio without addressing himself to the crucial question of how this alteration can be justified.

(3) But lastly, does reflective reasoning in the law incorporate the inductive policy of increasing confidence with positive instances and decreasing it with negative instances? There is, it seems to me, no short answer to this, and we must exercise care in approaching it. In the first place, it is clear that if reflective reasoning in the law does incorporate the inductive policy it cannot incorporate it in this form. Since with reflective reasoning we are not reasoning to the expectation that something will be the case, it is not strictly appropriate to speak of having confidence in the conclusions which reflective reasoning point to. But then, the inductive policy is actually very elastic and is not completely wedded to talk of confidence; changes in the description of the aim of induction can be allowed. For example, we could speak of positive or negative instances providing stronger or weaker grounds for thinking that our conclusions are correct. But if we modify the inductive policy in this way, then a definite answer to the question whether reflective reasoning adheres to this policy becomes considerably more difficult to provide.

We are pushed in two different directions. On the one hand, it does seem that the inductive policy as modified just misses the point of reflective reasoning in the law. For, the object of such reasoning is to determine as best as can be done what a positive or a negative

instance of a particular concept is, not to collect, search for or anticipate them. The point of reasoning about negligence, in other words, is not to make it possible to compile a list of cases which would, under some presupposed definition, count as negligent. Rather it is to determine what a case of negligence is like, and particularly, whether the case at hand is such a case. If positive instances of a presupposed definition is what a judge is after in his deliberations, then he will never lack for them. There is an indefinite number of them, all those real or conceivable cases which satisfy the criteria which are presupposed. But judges do not bring up cases for that purpose; the mere collecting of cases is not what they are interested in.

Still, on the other hand, if what I have been suggesting is correct and reflective reasoning in the law and elsewhere is reasoning from and to particulars, then one might expect that more and more cases which are clearly cases of negligence and similar to the instant case are bound to exert a pressure on our thinking, affecting the plausibility of our judgement about the instant case. If a judge is considering the issue of, say, whether the foreseeability of harm is a necessary ingredient of the concept of negligence, then if no case of near-negligent behaviour which lacks the ingredient of foreseeability is sufficiently similar to paradigm cases of negligent behaviour, then this fact, and the pressure of the cases which can be brought out, certainly should provide a stronger reason than that provided by a single case for thinking that foreseeability is a necessary ingredient. But now it appears as if reflective reasoning is at least guided by the spirit of the inductive policy.

Suppose we change the example and view the matter from a slightly

different perspective. A precedent has recently been set concerning the interrelated family of terms 'death', 'killing', and 'murder'. A Superior Court judge in Massachusetts has recently instructed the jury hearing a first-degree murder case to find the defendant guilty on the grounds of grievous assault (which was not contested) and evidence from doctors that although the victim's heartbeat and breathing were maintained for a considerable time after the assault by means of life-support systems, still from the time of the assault on there was a total cessation of brain wave activity which was directly linked to a blow on the head suffered in the assault. Now this case will undoubtedly be cited in future cases in support of the claim that so-called 'brain death' constitutes the single and sufficient condition for death. But suppose we engage in a bit of fiction and imagine that twenty years hence a judge is considering the question of whether in the case he is presiding over the fact of brain death should count as a sufficient condition for death. Suppose as well that for every previous case where brain death was a material fact of the case, brain death was ruled to count as death simpliciter. These cases, we can safely suppose, cover a wide range of legal issues, criminal and non-criminal, from murder to probate.

What we are interested in here is what sort of reasoning the judge would engage in to support the conclusion that, in this instant case, brain death is death. Consider these two sketches of arguments the judge might employ: (I) The Massachusetts case set the precedent that brain death constitutes death. Since then, every case of brain death has been ruled to be a case of death. Therefore, this case of brain death should be ruled to be death. (II) Since the Massachusetts case various judges have considered various instances where brain death

was a material element. Moreover, that particular material element has been found associated with a wide variety of other facts, in murder trials, and in probate trials. In each case the judge has ruled that there was sufficient reason to view the instant case as relevantly similar to the precedent to rule similarly, i.e., no countervailing feature of the case was viewed as counting as a reason against viewing brain death as death. Now it is possible to discern various analogies between these previous cases, including the precedent, and the instant case. Moreover, there are no significant disanalogies between this case and the others. Therefore, in the light of these cases and these analogies, this case of brain death should be ruled to be death.

The first thing to notice about (I) and (II) is that each in its own way exemplifies the common law rule of stare decisis: in both cases the argument appeals to previous decisions and previous arguments, and the decisions are viewed to be immune from retrospective revision: they are law. Moreover, both arguments satisfy the spirit of the modified inductive policy: previous decisions and previous arguments have an effect on the way the instant case is viewed. Yet, the first argument embodies and depends on an inductive argument whereas the second does not. With (I) what we have is the claim that the original Massachusetts case presented, in effect an inductive inference from the particular decision to the generalisation that in every case brain death is death. On the basis of that (I) concludes that the instant case should be ruled similarly, given the confirming instances which are on the books. The suggestion here is that the ratio decidendi of the Massachusetts case is a hypothesis which subsequent cases have tested, and, as it has happened, these tests have resulted in confirming evidence in support

of the ratio-hypothesis. (An account of judicial reasoning in inductive terms not substantially different from this one can be found in Cardozo, 22 ff.)

The second argument cannot be expanded in this inductive fashion. With (II) an entirely different approach is being taken. No hypothesis is in the background, no confirming instances are acknowledged. Moreover, there is with (II) no reliance on the projectibility of the key term 'brain death'; nor is the assumption made that all of the previous cases share some common set of features by means of which they could be grouped together as a sample set. Indeed, the multiplicity and variety of situations in which the material element of brain death has been found in previous cases is emphasised. Analogies and the absence of disanalogies are noted. Parallels are referred to and are treated as grounds for the appropriateness of treating the instant case like previous cases. But most significantly, these previous cases are considered not just in terms of the decision which was reached, but also in terms of the grounds for the decision in each case.

Now I would like to suggest that (I) is not at all the sort of reasoning which a judge ought to engage in when he is reflecting on the instant case. It seems clear enough that it would be infelicitous for him to reason in this way, and there is no doubt that if he were to attempt to justify his decision by citing something like (I) in his opinion, the case would be appealed and the appeal sustained. For (I) is a non sequitur in this context: it is conceivable that (I) might be employed by someone interested in predicting how various judges will decide cases involving brain death in the future (the sort of argument in the background of a claim like 'Given the long history of cases where

brain death was ruled as death, it is unlikely that any judge would want to upset to precedents that have been laid down by ruling otherwise'). But the judge himself cannot reason like that: he must look at previous opinions not just previous decisions. From the mere fact that judges in previous cases have ruled that brain death is death nothing whatsoever follows about how this case should be decided. But, from the fact that judges have ruled that brain death in various cases in the past, and from the particular features of those cases which led the judges to rule in the way they did, it may follow that the present case should be decided similarly. Hence an argument like (II) is at least pertinent to the judicial question. Moreover, it could provide the answer to the judicial question.

But, we are still left with the fact that (II), which sketches a reflective argument, seems to rely on the spirit of the inductive policy. There is with (II) an awareness of the weight which previous cases carry, despite the fact that (II) is not a piece of expectation reasoning, and is definitely not an inductive argument. Does reflective reasoning in the law incorporate the inductive policy or doesn't it?

Once again, there is no simple answer to this question. One thing does seem clear from our brain death example, however, and that is that the weight which previous cases have, and the pressure which they exert on deliberations about instant cases, is not accumulative or in any sense quantitatively measurable. And this is a feature of judicial reasoning which is very apparent. For often the weight which a single clear case carries and the pressure which it exerts is far greater than the weight carried and pressure exerted by several contentious or dubious cases. For inductive reasoning, of course, it is primarily the

mere number of positive instances which is important. (Primarily, but not exclusively, since obviously the same experiment performed over and over again does not provide more and more inductive support; scientists are also interested in the variety of test-conditions which yield confirming evidence for a hypothesis.) But further development of this issue must wait on a more extensive treatment of the nature of reflective reasoning itself. We have seen enough to realise that reflective reasoning is not a species of inductive reasoning, the differences are too great. But we have also reached the point where an extensive characterisation of reflective reasoning as a legitimate and useful form of reasoning is needed.

Footnotes

Chapter IV: Expectation Reasoning

(Page 168)

¹Judith Jarvis Thomson has denied this (Thomson (1)). She claims,

At all events, the question whether one thing is a reason for another is very often quite obviously a question for investigation rather than reflection. One thing that seems plain is that if 'This is S' does not imply 'This is P' it will always be at best a matter of fact, to be established by investigation, that the first is a reason for the second. (272.)

On one level we are not disagreeing: it being a matter of fact that something is a reason for something else is perfectly compatible with it being a question about the concept reason whether this is a reason for that. On another level, however we are disagreeing. It is, apparently, Thomson's view that factual questions can always be settled factually, by investigation and fact gathering. This is true when the factual question is expressed in terms of concepts which are (or are stipulated or assumed to be) determinate. But not all questions of fact are expressed in determinate terms, nor do we always assume them to be determinate. 'Is this carelessness?' is a factual question, but it is, given the non-determinate nature of the concept employed, a question which requires reflection over the facts. Facts about what was done, how it was done, under what circumstances and in what context it was done, are essential to the deliberation; but these facts do not always settle the issue. It is my contention, that reason is itself a non-determinate concept, a concept whose proper application in some cases requires reflection.

(Page 171)

²It might be objected at this point that what I have been blithely calling 'inductive reasoning in mathematics and logic' is really a kind of semi-deductive reasoning. The objection would probably go like this: Recursive procedures are not external to logical or mathematical calculi, rather they are legitimatised in set theory. In set theory there are theorems of mathematical induction which are provable schemas. That is, they are deductive consequences of the axioms of set theory. And, as such, their instances are justified in the normal, deductive manner. It is therefore misleading to talk about mathematical induction as if it were on all fours with empirical induction.

The point is well taken that there are clear and obvious differences between recursive arguments employed in logic and inductive arguments used in, say, physics. But, the mere fact that induction schemas are theorems of set theory does not mean that what the logician uses to prove his meta-theorems about infinite sets of wffs and what not is not inductive reasoning. These inductive schemas, one must realise, are derived from the Axiom of Infinity, the axiom which posits the existence of infinite sets by means of a recursive procedure: there is a set which is such

that (for example) 0 is in it, and for every natural number n , n 's successor is in it. Now this axiom is, in effect, a principle of induction, for it establishes what procedure one must follow in order to determine whether a set is infinite. This procedure is, of course, an inductive procedure; it couldn't be anything else. So, the logician escapes the problem of justifying mathematical induction in a way which those who are troubled with the problem of justifying all induction would, generally speaking, heartily approve of: deducing particular instances of induction from a principle of induction.

(Page 175)

³M. R. Ayers has expressed the view that there are no special problems about verifying or analysing either counterfactual or subjunctive conditionals:

What is it about a conditional statement which makes it 'counterfactual' or 'unfulfilled'? Certainly not any feature of the statement itself or its expression. Whether a condition is fulfilled is like the question whether or not a factual statement is true: the answer depends on something quite outside the meaning, form or category of the statement or the sentence by which it is expressed. (349.)

I take Ayers's point in this article to be substantially the same as the point I have been trying to make in the text with regard to inductive arguments: there are questions about the logical status of inferences and there are questions about the acceptability of the argument--but these are different questions.

(Page 188)

⁴It should be noted that some Realists, following comments made by O. W. Holmes (in "The Path of the Law") seem to be saying just this. For Holmes the law is nothing more than a set of predictions that if a man does or omits certain things he will be made to suffer the consequences (458). As it has been repeatedly pointed out, however, if Holmes's characterisation of the law is appropriate from the perspective of the litigant, it is wholly inappropriate from the perspective of the judge. The judge's reflection is directed to a decision not a prediction.

CHAPTER V

REFLECTIVE REASONING

Bringing the Pieces Together

In the past three chapters I have been interested in the nature of deduction and induction. The points raised and pursued in these discussions have centred around the general question of how deduction and induction satisfy two proposed requirements of reasoning, namely the requirements of utility and security, epistemic advance and legitimacy. In the course of these discussions a picture of the nature of reasoning itself has begun to emerge and in this chapter I want to expand upon and clarify that picture. It is first necessary, however, to see how far we have already come by isolating and scrutinising some of the points which form the outline of the picture I have in mind. What we need to do is to bring the pieces together. Suppose we begin with a major piece of the puzzle.

A key result of our investigations concerns what might be called the priority of the informal example. It was argued that if we pursue the question of how, in the case of deduction, the requirement of legitimacy is satisfied, we are eventually forced to move from a consideration of particular and informal examples of valid deductive arguments. The priority of the informal deductive argument is manifested by the manner in which the formal logician is forced to argue when the primitive bases of his logic--the axioms and rules--are questioned with respect to their legitimacy. And a parallel point was made in the case

of induction: worries about the legitimacy of inductive arguments are in no way removed by the systematisation of inductive logic or the introduction of the quantitative machinery of probability theory.

My concern in these discussions was to discover the source of the legitimacy of formal deductive or inductive procedures, not to draw attention to any hidden faults in them. Deductive proof from axioms and in accordance with rules of inference is perfectly legitimate and perfectly conclusive proof. Moreover, we have every reason to think that the inductive policy formalised by the principle of induction is the correct policy of expectation reasoning. These facts were not in dispute. Indeed, if anything they were reaffirmed by our treatment of the source of legitimacy. For if one is wedded to the notion that questions about the legitimacy of deductive or inductive procedures reach to the axioms, or first principles, and no further, then one is vulnerable to the sceptic who poses the question, What about these principles, why are they legitimate? As we saw in the earlier discussion of Deductivism, one soon runs out of first principles. But one never runs out of instances of these principles--the clear cases which can be multiplied indefinitely--and it is these which constitute the source of the legitimacy of these principles.

Now the priority of the informal example in the case of deduction is easily lost sight of since deductive logic is taught and learned as a body of rules and principles, syllogistic patterns, rules of inference and necessarily true propositional patterns. We tend to forget that most of the key rules and principles are obtained by generalising from patently valid informal arguments and necessary propositions expressed, not in the symbols of some deductive language, but in ordinary, natural

languages. And the same is true in the case of probability axioms and rules of inductive support. These have arisen out of quite ordinary inductive arguments; indeed the principle of induction itself--perhaps the most vulnerable logical first principle--rests on our countless inductive beliefs about the continuity and regularity of our environment.

But the priority of the informal example is easily lost sight of for another reason. As we saw, once a logic has been fully formalised it can be manipulated to generate an infinite number of necessarily true propositional patterns, or theorems. Deductive calculi, those paradigms of optimally determinate systems, have a life of their own. The axioms and rules enable us to reason from principle to principle without having to proceed from a consideration of particular instances. Thus, it is very plausible to suppose that, in the proper conceptual environment, reasoning from and to principles is a distinct and autonomous rational activity. Informal deductive arguments are left behind and their essence, their logical form, is retained and systematised.

But, if the power and perspicuity of formal deductive reasoning is (as I think it is) undeniable, it is nonetheless deceptive. For one may be lead to the view that the deductive logic machine is reasoning par excellence, autonomous and self-validating. This last supposition is contained in what I have labelled the 'formalist assumption': provability and validity--and so, legitimacy--is a function of the foundation of the closed, determinate system itself. But if deductive legitimacy is to be traced to, and is to stop at, the axioms and rules, then we are in the Deductivist quagmire once again. The remedy lies in reminding ourselves of the priority of the informal deductive argument. But this remedy leads to a dilemma.

This was the dilemma posed for the formal logician in Chapter II. On the one hand, the source of the legitimacy of formal deductive inferences is not the axioms or rules but particular and informal examples of these which strike us as patently valid. But on the other hand, the formalist assumption which locates the source of formal validity in the axioms and rules is motivated by a very persuasive consideration: reasoning within a deductive calculus provides genuine epistemic advance, and it does so without any recourse to informal deductive reasoning. But then, if the formal logician maintains the formalist assumption in the face of the priority of the informal deductive argument he is in effect disassociating formal logic from the legitimacy requirement, and formal logic ceases to be a system of reasoning. If he gives up the formalist assumption, however, he is thereby rejecting the very raison d'être of formal logic, since it is the formal structure of deductive calculi, the axioms and rules, which allows for the power and perspicuity of the reasoning carried on within such calculi.

It was while attempting to resolve this dilemma that two points of considerable importance to the larger issue of the nature of reasoning were raised. It was first noted that the formalist assumption expressed in terms of formal deductive logic is symptomatic of a far more general epistemological disease--Euclid's disease, or what Wittgenstein calls the "craving for generality" which is accompanied by a "contemptuous attitude towards the particular case". What we hit upon was in fact a substantial epistemological dogma which the later Wittgenstein and Wisdom tried to dislodge from the minds of philosophers. As applied to the question of the nature of reasoning, the dogma states that reasoning from principles (from axioms, rules, formulas, idealisations and general

statements of various sorts) provides legitimate epistemic advance whereas reasoning from the particular case to the particular case does not. The disease was also shown to be related to a craving for definitions and for necessary and sufficient conditions for the application of all concepts, with the attendant assumption that we do not really know what a term means or what a concept applies to unless and until we have definitions and sets of necessary and sufficient conditions. But then, it became apparent that the priority of the informal example was just one aspect of a far more general cure: the realisation that reasoning from the particular to the particular, from examples and parallel cases is prior to, and the source of the legitimacy of, reasoning from principles. The particular gives life to the general, and all reasoning ultimately resolves to a consideration of cases--Wisdom's 'mother's method'.

The second point also concerned the craving for generality, but this time the non-obsessional aspect of that craving was considered. What I was interested in was the raison d'être of formal systems. It was pointed out that reasoning to and from principles, reasoning carried out within and subject to the exceptionless rules of deductive calculi, is valuable reasoning because genuine epistemic advance is often the result. The new understanding which can result is an understanding of the calculus itself, and this typically takes the form of new theorems. And the same can be said, with some obvious modifications, for other formal systems which are organised and structured by deductive logic, namely those so-called 'interpreted' logical theories which axiomatise mathematical theories and various branches of the sciences. The epistemic advance which arises in these systems is considerable and of no mean

importance.

The lure of the formal system is its power and elegance, the ease with which we can move from point to point in the network of logically connected propositions and the simplicity of the formal proofs which can be constructed. These advantages are the result of the overall structure of formal systems, and this overall structure exemplifies a systematic property which I have called determinateness. Determinate systems possess the structural characteristics of an axiomatised theory, of the algorithm. And deductive calculi provide the model, the exemplar of system determinateness. But the key to system determinateness is located at the conceptual level: the logical coherence of a formal system is a direct function of the unambiguous logical connections of a deductive nature which exist between the concepts employed in the system. These determinate concepts are in turn governed by strict necessary and sufficient conditions which eliminate ambiguity and borderline cases. For the most part, determinate concepts are theoretical concepts, concepts embedded in a particular theory and explicitly defined in terms of other concepts of that theory. And there are examples of determinate concepts in ordinary contexts as well. Yet, in the case of deductive calculi, concept determinateness is manifested at the most basic level: the truth-functional interpretation of the key logical constants, connectives and operators.

By bringing these two points together the dilemma for the formal logician was resolved. The raison d'être of formal systems can be preserved: reasoning within these systems is powerful and elegant, and these virtues can be credited to the determinateness which these systems have acquired through formalisation. But for all its power, formal

reasoning is neither autonomous nor self-justifying, for ultimately formal procedures rely on informal examples of reasoning for their legitimacy. Thus, the formalist assumption, despite its initial plausibility and attractiveness, misconstrues the nature of reasoning. Formal deduction constitutes an idealisation of informal reasoning. The roots of formal procedures are permanently fixed in informal deductive reasoning from which they have been abstracted.

So, our search for the source of the legitimacy of formal deduction lead us to the level of informal deductive reasoning. But it then became apparent that our search had not ended. It was at this point that Mill's insight that, in the end, all reasoning is from the particular to the particular in a non-deductive fashion acquired a new importance. Three clues were then set out, each of which suggested that reasoning itself resolves to non-deductive, case-by-case reasoning.

The first of these clues came from a reappraisal of the path we had travelled. I had argued that when the question of the legitimacy of a particular deductive inference in a formal system was considered and when we followed the deductive model of justification, we were eventually lead to the axioms and rules of the system. But at that point deduction ceased to be an effective way of reasoning: we were in the position of Achilles trying to satisfy the Tortoise's unsatisfiable request for more deductive justification. But if deduction was no longer available to us, nonetheless reasoning was not hobbled since we could go on to consider instances of the axioms and rules whose necessity and validity are manifest. The question was one of justification, and it was obvious that if justification always and only consists in an appeal to general principles, then at some point

justification must end at some extremely general propositions which themselves cannot be justified in that manner. Yet, reasoning does not end here, for we can go on to consider particular instances of these general propositions.

The second clue came from a consideration of ranges of questions and problems which require reasoning but for which neither formal nor informal deductive reasoning seemed to be the suitable rational vehicle. Here we took a look at the law and the range of questions which are daily decided in courts of law. Two traditional jurisprudential theories were discussed in connection with the issues of the nature of the legal system and the character of judicial reasoning. Two conflicting views emerged. On the one hand the Formalist pictures the legal system on the model of a determinate or near-determinate system and on the basis of this model concludes that judicial reasoning is deductive in nature. The Realist, on the other hand, sees legal systems as mere collections of sociologically significant facts, in particular, judicial decisions and the laws on the books. There is, for the Realist, no clear sense in which judicial decisions are reasoned at all; rather these decisions are either the product of socially-developed intuitions, more or less arbitrary decisions about the meanings of key legal terms, or conscious attempts on the judge's part to reflect in his decision the majority view of his society. Both views, despite their differences, were argued to be founded on the same assumption about the nature of reasoning. That assumption is that if there is to be reasoning in the law it must be deductive in nature: either the judge reasons to a decision by following strict rules for the application of statutes and precedents to the instant case, or else he proceeds in a fundamentally non-rational manner

following his intuitions or responding to the will of his society.

Once again, this assumption forms part of a particular line of argument which Wisdom has repeatedly attacked. The argument runs like this. If we have a question which cannot be definitively answered by means of deductive reasoning, and if no more facts could be revealed which would be relevant to it, then the question ceases to be a genuine one and is a mere matter of words (vide Wisdom (3)). Now we often have at the end of the proceedings of a court the following situation: (a) there are no deductive relations to be found between precedents or statutes and the instant case; and (b) all the factual information which could have been brought out has been. Nonetheless, the question of whether the instant case was a case of, say, negligence or not may remain open. It is still a genuine question, and it is a question which requires further reflection. There are legal philosophers who have avoided the pitfalls of both Formalist and Realist camps and have argued that reasoning in such cases is possible. And these philosophers bring out the same point, that the reasoning, the reflection, that is required in the law consists of a consideration of parallel cases.

I suggested in this discussion that there is a natural temptation to think that judicial reasoning is exotic or extraordinary since there are strict procedural rules which define proper procedure in the courtroom. But reasoning from case to case, by comparison and contrast, is not some sort of 'alternative' logic the application of which is restricted to the legal forum or the judge's chambers? As we saw, the sorts of questions upon which legal judgements often hang are questions about the application of very ordinary notions indeed, notions like carelessness, risk-taking, foreseeing harm, and others. And even when a specialised

legal vocabulary is employed, the issues which require reasoning are issues which can be broken down into quite ordinary, and not particularly exotic, questions about everyday situations. What is significant about judicial reasoning, however, is that it is the working logic of the court. That is, the applicability of deductive procedures wait upon the answers to questions which require case-by-case reflection. Deductive inferences form part of the reasoning which leads to a judicial decision, but at the stage at which deductive reasoning becomes applicable, all of the perplexity of the instant case has been removed.

The last consideration suggesting that reasoning resolves to case-by-case reasoning was uncovered by looking more closely at the questions which are at issue in the courtroom. It became apparent that these questions are questions about legally relevant concepts. Typically, the judge in adjudicating a case is faced with the problem of putting the facts which are open to view into a pattern which will enable him to see the case as one of a certain kind. The judge is faced with the problem of finding the correct description of the situation, act or thing before him; and this is the problem of seeing it for the kind of situation, act or thing it is. The judge, that is to say, is dealing with questions about concepts. And these questions are not restricted to the legal sphere. Examples abound of questions of a similar sort elsewhere, in other fields, ordinary and specialised. These are questions which give rise to puzzlement and perplexity. They are questions about borderline cases, examples of things which are not definitely of one sort rather than another. Moreover, it was suggested that the incompatibility which exists in moral conflicts can best be understood as a dispute over just this sort of conceptual question.

But when we concentrate on questions about concepts, when we consider examples of problems requiring reasoning which turn on questions about concepts, we notice that deductive reasoning is usually ineffectual. For in order for the power of a deductive inference from premisses to conclusion to be at our disposal we must be able to agree on the premisses. But deductive premisses--whether they are principle-like or descriptions of particulars--employ predicates whose applicability is fixed. They rely on, in other words, concepts which are, or are assumed for the sake of the inference to be, determinate. Deductive premisses, therefore, presuppose that questions about concepts have been answered. But for the sorts of cases we considered--in the law, in morality and elsewhere--it is questions about concepts which are the source of the perplexity. To answer questions about the correct description of an act or thing from a legal, moral or aesthetic point of view we do not always have at our disposal necessary and sufficient conditions for the application of the appropriate concepts. And when we do not have these rules, when the issue is whether these rules apply in the case in question, then deductive reasoning has no foothold, it cannot even begin. What is required is a more fundamental way of reasoning, a way of reasoning which provides reasons for thinking that something is or is not of a certain kind, is or is not properly described in a certain way. But to see that deductive reasoning presupposes much of what is at issue with questions about concepts is to see that there is a need for reasoning at this crucially important conceptual level. We should, as a consequence, be wary of the view that deduction is the paradigm of reasoning; as Wisdom warns: "We fasten on the excellence of the deductive step from premisses to conclusion, neglecting the steps required for establishing the premisses themselves" (Wisdom (3)).

Now the excellence of the deductive step is a consequence of the logical grammar of the connectives 'if...then___', 'and', 'or', and 'not' and the operators 'all' and 'some' (and variants). And although there is some room for doubt concerning the adequacy of the truth-functional interpretation of these particles, once we agree to that interpretation, the rules of the deductive inference game have been set and there is no room for doubt concerning what deductively follows from what. But, significantly, both arguments for and arguments against the truth-functional interpretation of these particles tend to bring up examples of their use in ordinary contexts which either conform to or are at variance with the truth-functional interpretation. So, here again, the reasoning which is required to explore the boundaries of the grammar of certain terms is reasoning from example. In the end, the reasons we have for acknowledging the excellence of deductive moves are drawn out of our reasoning about particular examples. The very possibility of deductive reasoning depends on reasoning which is not deductive.

These clues, then, point to the conclusion that there is more to reasoning than simply employing deductive procedures of proof. They also suggest that reasoning from examples and parallel cases, or what I have called reflective reasoning, is a more fundamental way of reasoning. For the path we have travelled in our discussion of deductive reasoning has lead us directly to reasoning from case to case: As formal deductive logic relies on the legitimacy of informal deductive arguments, so these arguments rest on the logical grammar of a handful of connectives and operators. But that the logical grammar of these particles is what it is is a consequence of our recognition that these particles typically play the truth-functional roles which deduction requires them to play.

This recognition is founded on countless examples of propositions employing these particles which could be brought out in support of the claim--were it to be challenged--that these particles do play these logical roles. That we can legitimately infer a conjunct from a conjunction is obvious--that is part of what 'and' means. But the legitimacy of this move is based on our understanding of the logical grammar of 'and'. This understanding, however, can only be rationally defended by a consideration of case after case where 'and' functions in this way and not in any other way.

It is informative, and quite natural, to think of deduction as a game the moves of which are governed by rules. Some of these rules are constitutive, e.g. rules of well-formedness, rules of sentencehood. Such rules are conventional and, in a sense arbitrary: as long as they can be unambiguously applied and do not lead to inconsistency, formation rules can take any form. But deduction is a game with a point. The inference rules of the deduction-game are supposed to capture what rationally follows from what. But whenever rules have a point, their success depends on the success of particular instances of them, those instances of which the rules are generalisations, summaries. In the case of deduction, the rules of the game are summaries of countless examples of the logical grammar of the connectives and operators which are drawn from ordinary discourse (and, it might be supposed, common sense). Thus, whenever we use a deductive rule to identify some propositional, argumentative move as a case of what rationally follows from what, we are doing something which is formally different, but not rationally different from bringing to bear countless analogous moves which isolate the reasons why that propositional move is a case of what rationally follows from what. Deductive procedures are thus examples of

reasoning which are not only based on, but in the end resolve to, countless instances of reasoning from case to case.

It is important that we compare the path we have travelled with that travelled by Mill. It will be recalled that Mill's treatment of deduction was directed to the question of how deductive inferences satisfy the epistemic advance requirement whereas I have been concerned with the source of the legitimacy of these inferences. Our paths were joined at one important point: reasoning from principles ultimately resolves to reasoning from particular to particular. In Mill's case, however, the realisation that principles are no more than summaries of particular cases, coupled with the misleading metaphor of the conclusion of a deductive inference 'being contained in' the premisses, lead him to search for the genuine inference, in the background, which provides epistemic advance. I, on the other hand, have acknowledged what in any event is an undisputable datum in any discussion of deduction, that epistemic advance can result from deductive reasoning. I went on to try to account for this by centering the discussion on the question of the source of the legitimacy of deductive reasoning. Our paths were joined, but then they sharply diverged. Mill's search for the genuine inferential move, and his preoccupation with propositions about matters of fact, lead him to the conclusion that it is induction which supplies the epistemic advance in all cases. On my view, the fundamental form of reasoning is, as Mill saw, reasoning from case to case; but it is not always inductive reasoning from case to case.

At this stage it was necessary to turn our attention to inductive reasoning. It was argued that, as with deduction, if we adopt the tactic of seeking the source of the legitimacy of formal inductive procedures--the

logic of evidential support--we are lead to a consideration of the legitimacy of formal inductive reasoning. But to seek the source of the legitimacy of inductive reasoning is just to attempt to solve the perennial problem of justifying induction. By applying some of the lessons we had learned, we noted that this problem could not be solved by the favoured method of appealing to a principle of induction, since whatever justification there is for such a principle amounts to the justification there is for the instances of it which the principle is supposed to legitimatise. What is required is an understanding of the kind of reason an inductive reason is, and, significantly, that understanding arises out of a consideration of particular examples of inductive reasons. So, here again, the source of the legitimacy of a rational procedure is found in the results of reasoning which is more fundamental.

But in light of Mill's claims about induction, and the analogy between inductive reasoning and reflective reasoning, I went on to try to distinguish between the paradigm case of an inductive inference and reasoning from case to case. I kept the discussion general and tried to capture the conceptual requirements of inductive reasoning. The result was several clues which pointed to the conclusion that the sort of reasoning the judge employs to solve legal questions does not fit the pattern of inductive reasoning. Inductive reasoning, it was argued, is expectation reasoning and as such differs in orientation from reflection, reasoning which is directed to questions about concepts. Inductive reasoning is paradigmatically orientated to the prediction, to the expectation that further investigation will disclose that the regularity recorded in the premisses was correctly applied to the induced individual.

As such, inductive reasoning presupposes decisions about the determinateness of the concepts which identify the sample set as well as the concepts which identify the regularity projected onto the induced individuals. Induction, therefore, is reasoning with concepts, but not always reasoning about them.

There is a significant analogy between inductive reasoning and reflective reasoning from case to case. An inductive reason is a reason because of the kinds of particular cases which it brings to bear on the inductive conclusion. Induction is reasoning which utilises analogies and directs our attention to the evidential weight of similar cases. Indeed, the inductive policy itself--the policy of increasing confidence with positive and decreasing it with negative instances--is a policy which is often embodied in reflective reasoning. The differences in orientation suffice to cast doubt on the suggestion that reasoning by parallel cases in order to investigate in an a priori manner the boundaries of concepts is inductive reasoning. However, the similarities between inductive reasoning and reflection do seem to suggest that the former is a species of the latter, that inductive reasoning is reflection restricted to a determinate conceptual environment. Inductive reasoning is thus reflective reasoning constrained by its orientation, the role it plays as expectation reasoning.

We are now at the point where a more complete picture of what I want to claim is fundamental to the nature of reasoning can be developed. What I want to develop is a picture of reflective reasoning from case to case by means of an extended discussion of the question of the legitimacy of reasoning in this way. Given these last comments, the most appropriate approach to this project is via a treatment of analogical

reasoning, instances of which form the borderline between inductive and reflective reasoning.

Analogical Reasoning

Analogical reasoning is reasoning by example, from particular case to particular case. Specifically, it is reasoning based on a proffered analogy between two or more things--where 'things' in this context must be given the widest possible sense. Things are said to be analogous when they are in a relevant respect similar or dissimilar and one reasons analogically when one cites similarity (or dissimilarity) in one or more respects as a reason for treating two or more things as similar (or dissimilar) in other respects. But since there are various, perhaps countless, respects in which two things can be relevantly similar, there are various kinds of analogies. Arguments from analogy are consequently a diverse lot differing as to the nature, or one might even say the personality, of the proffered analogy.

There is, however, a tendency among those who talk about argument from analogy to work with a cramped notion of an analogy. Scorning the literary analogy, the allegory and the parable as so much purple prose, these authors opt for the formal logician's favoured sense of similar: two things are similar if, and to the degree that, they share properties. The result is a set-theoretical understanding of an analogue. Thus, one textbook defines an analogical argument as

an argument based on the inference that if one individual or class shares one or more properties with a second individual or class, one or more other properties of the first are also shared by the second. (R. Olson, 384.)

Fallacious analogical argument--the 'fallacy of faulty analogy'--becomes on this view the mistake of predicting on the basis of shared properties

which are trivial, irrelevant or not sufficiently characteristic of the kind of thing one is talking about. In short, the view is not only that analogues are only identifiable by means of the properties they have in common, but also that analogues are closer, or more stable for predictive purposes, the more properties they have in common. By implication, then, the limiting case of analogy is identity and the ideal analogical argument is the deductive argument one gets by replacing the premiss which cites an analogy between two things with the premiss which cites an identity between them.

Now here is a case where, by insisting on cramping a notion, one is, paradoxically, also mercilessly stretching it and related notions. Here the victim is the notion of similarity. For on this account, one way (albeit an ideal or limiting way) in which two things can be similar (and analogous) is if they are not two things at all, but one. But this surely won't do: although two things can be the same only if they are not different, two things cannot be similar unless they are different. Still, if this objection should lead us to think that this notion of an analogy is not even part of the story, then I think we will have been misled. For some analogues are things which share properties, and some arguments from analogy are inferences from the fact of property-sharing to the prediction of further property-sharing. Such is inductive analogical argument.

Analogical arguments based on this property-sharing conception of an analogy are inductive arguments from particular to particular (instances of the sort of inductive argument I called type (B) in the previous chapter). The fact that analogical arguments are arguments from a particular case to another particular case has caused some difficulty among those who, following Aristotle (Analytica Priora, II, 23), have

characterised induction as a process of reasoning from (some set of) particulars to a generalisation (and then, perhaps, to a prediction). We have already argued above that this is an unnecessarily restrictive conception of induction, but for further evidence, consider these two arguments:

- (1) Our experience with several cases of Pernicious Fever has been that whenever the patient reaches the stage where the skin becomes pallid, death by kidney failure follows within 24 hours. Mr. Brown has reached that stage; so it's high time we find a kidney machine.
- (2) We have on record the case of Mr. Smith who, after being exposed to DDT in the storeroom of A & B Garden Supplies developed a high fever, dizzy spells, cold sweats and loss of appetite which lasted for a week whereupon his skin became pallid and he died of kidney failure the next morning. Mr. Brown has been exposed to DDT in that same storeroom, has had the same symptoms for a week and has now begun to look very pale indeed. So it's high time we find a kidney machine.

Now both of these arguments are straightforwardly inductive, they are both arguments which lead to a prediction about a particular case which is based on previous experience. The difference between them is that whereas (1) bases a prediction on several previous cases, (2) bases its prediction on the fact that a single previous case and the present one share properties, the cause and symptoms of a disease. There is little doubt that (1) is an inductively stronger argument--although (2) remains a justifiable reason for action. Moreover, (2) shows the appeal of the property-sharing conception of an analogy: there is something to the claim that the more properties things share the more reasonable it is to suspect that they share other properties.

As a form of induction, inductive analogical argument presupposes conceptual determinateness. The properties shared, and the properties which the argument predicts the analogues will share must be assumed to be, for the purposes of the inference, unambiguously possessed by the analogues. And in a case like our Pernicious Fever example, this is a natural enough assumption to make. Yet, inductive analogical reasoning is generally regarded as a very insecure form of reasoning, and for good reason. If we want to argue on the basis of an analogy, a similarity between two things, our concern is not with the mere number of properties shared, but with the number and relevance of the properties shared. (2) is a defensible argument because Mr. Smith and Mr. Brown share symptoms; sharing the property of possessing a large stamp collection or having read Finnegan's Wake before reaching the age of 22 does not (unless a suitable story is told) improve the analogy, even if it does increase the set of shared properties. For this reason inductive analogical reasoning is viewed with some caution: fallacious and bizarre inductive analogical arguments are formally indistinguishable from sound and reasonable ones.

But then, judgement is required to separate the fallacious from the reasonable inductive argument from analogy. But what are we looking for? Presumably something more than the mere sharing of properties; we are looking for relevant similarity. If we were wedded to this conception of analogical reasoning we would have reason to be sceptical about its value. There is something suspicious about counting shared properties in order to ascertain analogies:

There is a case on record of a woman who, having had all she could stand of her husband's late-night carousing, waited for him to stumble home, sink into his favourite chair and (she thought) doze off before she

shot him several times in the chest. An autopsy revealed that, rather than having dozed off, he had died of a heart-attack and had been dead a full half-hour before she shot him. (She was convicted of attempted murder.) Now suppose we set this case alongside another one which is identical in all respects to the first except that the husband was alive at the time of the shooting. Are these two cases analogous? Do we have a close analogy here? On the property-sharing account they are about as ideally similar as any two, non-identical, cases are going to get. But the single property they do not share suffices to block the inference either to 'They are both cases of murder' or to 'They are both cases of attempted murder'.

Or we might think here of the often cited judgement given by B. Cardozo J. in the case of Hynes v. New York Central Railway Co. ((1921) North Eastern Reporter, 131, 898; and see also R. Cross's discussion of this case in his Precedent and English Law, 187, ff.). In this case the young Hynes was electrocuted by a falling high tension wire while he was trespassing on the railroad's property. Hynes had been swimming in the Hudson River and had climbed up on a diving-board which was built onto the bulkhead which belonged to N. Y. Central but which extended over the river, which was public. He was diving off the board when he was struck and killed by a falling wire from one of the defendant's poles in consequence of the admitted negligence of the defendant's maintenance crew. The railroad argued that Hynes was a trespasser and as no duty of care is owed to trespassers, they were not liable in tort. Their argument, in fact, was that the instant case shared more features with previous cases concerning landowners' duties to trespassers than it did to previous cases concerning landowners'

duties toward those exercising public rights (e.g. the right to swim in public waters). Justice Cardozo threw the railroad's case out and ruled for the plaintiff. The Justice, while admitting that Hynes was technically a trespasser, charged the railroad of pushing the analogy to a "drily logical extreme" while completely overlooking the consequences of such a property-sharing analogy. Cardozo argued that by forcing the analogy between the instant case and duty-to-trespasser cases we are being asked to concentrate our gaze on the diving-board and to focus on the fact that that diving-board is private property while ignoring the public ownership of the water and the space of air surrounding it. This puts the circumstances in a distorted light. We miss what is relevant to the case because the analogy is directing us toward what is irrelevant. In support of his judgement, Cardozo drew attention to the consequences of the railroad's analogy: duties to care arise or are extinguished in different zones or strata:

Two boys walking in the country or swimming in a river stop to rest for a moment along the side of the road or the margin of the stream. One of them throws himself beneath the overhanging branches of a tree. Another perches himself on a bough a foot or so above the ground.... Both are killed by falling wires. The defendant would have us say that there is a remedy for the representatives of one, none for the representatives of the other. (Ibid., 900.)

Thus, as far as Cardozo J. was concerned, reasoning from analogy can lead to mere legal hairsplitting and he was determined not to be misled by it.

Cardozo's is an example of one of three prevailing attitudes towards argument from analogy. Cardozo would have us avoid such arguments-- the risks we are taking in relying on them are just too great. Another attitude (the Formalist's rejoinder) is that, whether it offends our

intuitions or not, the only precise way we can make sense out of analogues is in terms of property-sharing; any other conception of an analogy just muddles matters. The last prevailing attitude is that, in the end, an analogy is whatever we think it is; we must simply develop an intuitive feeling for such things. Each of these attitudes is unnecessarily pessimistic. What we need to do is learn from the cases and try to uncram the notion of an analogy.

There is little doubt that the value of an inductive analogical argument is a function of a context. One has a point in mind when one draws attention to the properties which two things share. In the Pernicious Fever example, (2), there is a background which is presupposed without which that argument would lose its point. Drawing attention to symptoms only makes sense in the context of some common beliefs about the nature of disease. And in that context there is a harmony of opinion concerning what, for example, is to count as a high temperature or a loss of appetite. It is significant that, within this general context of disease and treatment, there is a range of degrees of conceptual determinateness, a range which has an effect on the utility of the resulting arguments. The conceptual connections between symptoms and nature of disease, and between nature of disease and nature of treatment form the logical skeleton of the context in terms of which the analogy between Brown and Smith makes sense. In this context, the value of the argument becomes a question not only of the extent of symptom-sharing, but also of the definiteness of the symptoms. There is a range of temperatures which could in various circumstances count as high or a sign of ill-health, but the chances that Brown has the same disease Smith died of are increased still further if both had higher temperatures at night than during the day, if both failed to respond to febrifugal drugs, and so

forth. The rationale of identifying a closer analogy with more relevant properties shared manifests itself. But these analogies, and arguments based on them, presuppose a stable conceptual environment.

Compare our Pernicious Fever argument with one the point of which is to determine what it means for two people to be suffering from the same disease. In this latter argument, symptom-sharing is bound to be a key issue. But notice the different role it plays. The question is not the empirical one of whether two people as a matter of fact share symptoms. It is rather the conceptual one of whether the fact that they do counts in favour of saying that they are suffering from the same disease. The question is whether symptom-sharing is relevant to the judgement that two people have the same disease, and if relevant, whether it is a good or sufficient ground for that judgement. It is, moreover, a question which does not require more observation, pulse-taking or thermometer-reading; it requires reflection not prediction: Brown's death by kidney failure, in this argument, would have to be argued to be relevant to the question of whether he had what Smith had.

I want to say, then, that the property-sharing conception of an analogy has a significant and legitimate logical role to play, but that it can be employed only after other questions have been resolved. We have encountered these sorts of questions before, they are questions about concepts. To see how argument from analogy can be directed toward such questions is to begin to release the notion of an analogy from the formal logician's stranglehold.

Let's consider some examples (the first is a summary of the majority opinion in Haseldine v. Daw discussed above):

(3) A manufacturer of products, which he sells in a form indicating

that he intends them to reach the consumer in the form in which they left him with no reasonable possibility of intermediate examination, owes a duty to the consumer to take reasonable care that the products will not result in an injury to the consumer. In this case the plaintiff sustained injuries as the result of the collapse of an elevator which was caused by the negligence of a firm of engineers employed to repair it. Given that there was no reasonable possibility of intermediate examination, the repairer stands in the same relation to the elevator as the manufacturer to his products. Therefore, the repairer owes a duty of care to users of the elevator.

(4) "For as the eyes of bats are to the blaze of day, so is the reason in our soul to the things which are by nature most evident of all."

(Aristotle, Metaphysics, II, 993.)

(5) When Jones prevented that man from jumping out of the window to his certain death it might have seemed that he was just meddling in someone else's affairs, but really he was a Good Samaritan.

(6) Doesn't a pregnant woman have a special responsibility toward the unborn child she is carrying, a responsibility which would make having an abortion immoral even in the case of pregnancy due to rape?--

...suppose it were like this: people-seeds drift about in the air like pollen, and if you open your windows, one may drift in and take root in your carpets or upholstery. You don't want children, so you fix up your windows with fine mesh screens, the very best you can buy. As can happen, however, and on very, very rare occasions does happen, one of the screens is defective; and a seed drifts in and takes root. Does the person-plant who now develops have a right to the use of your house? Surely not--despite the fact that you voluntarily opened your window, you knowingly kept carpets and upholstered furniture, and you knew that screens were sometimes defective. (Thomson, (2) 132.)

It would be very difficult, without risking distortion or oversimplification, to say anything of a general nature about all of these examples. Analogies, as I have already said, are a mixed bag. One thing we can safely do is to compare these examples with the inductive analogy in (2). What we see is that with (3)-(6) there is a different notion of similarity at work: With (3) the analogy is a parallel relationship which is linked by means of a common feature, the absence of a reasonable possibility of intermediate examination. We are being asked to see the relationship between repairer and chattel repaired in the same way as the relationship between a manufacturer and his products. The analogy is a similarity of relationship which is reflected indirectly as a similarity of relata. The repairer and the manufacturer are analogous, for the purposes of this argument, solely by virtue of the similar relationships they bear to products and repaired chattel. The logical relationship here is brought out clearly in the case of Aristotle's analogy where the time-honoured pattern A is to B as C is to D is exemplified. In both of these cases we have the sort of analogy which Kant was to characterise as "not...an imperfect similarity of two things, but a perfect similarity of relations between two quite dissimilar things" (Kant (2), 106; Kant gives as an example of such an analogy the following: "I never can do anything to another man without giving him a right to do the same to me on the same conditions; just as no mass can act with its moving forces on another mass without thereby occasioning the other to react equally against it" (Ibid., fn. 1)).

With (5) and (6), however, we see a more complete departure from the property-sharing conception of similarity and analogy. Lawrence Becker in his "Analogy in Legal Reasoning" has suggested that these sorts

of analogies are "dynamic", they are analogies for which the ideal is not identity but complete intertranslatability (251-252). In the case of (5), the argument turns on seeing Jones's behaviour in the light of the behaviour of the Good Samaritan of Luke 10:30-35 (and, as well, of other Good Samaritans, other analogues of the original helper of the man who fell among thieves). What we are being asked to consider is the affinity between what Jones did and what the Good Samaritan did, to translate Jones's behaviour in the 'language' of the Good Samaritan. The fact that Jones did not bind any wounds, pour oil and wine, or pay any inn bills is quite besides the point. If there is something wrong with the analogy, that isn't it. The Good Samaritan's behaviour is conceptually related to, by being a paradigm instance of, doing good by one's fellow man in times of need; and the point of the analogy in (5) is to read Jones's behaviour as another instance of that sort of behaviour. It is as if we are being asked to try to see Jones in the clothes of the Good Samaritan--rather than in the clothes of the meddler--and to be impressed by the close fit.

But what are we to say about J. J. Thomson's analogy? Here we have a bit of science fiction which is supposed to help us answer the question of whether a pregnant woman has a moral responsibility toward her unborn child. And there is no doubt that it does help us in this regard--the question is how it does. What is going on here is something like this. The question Thomson is addressing herself to is asked in such a way as to bring one feature of the issue of abortion into prominence--the mother's responsibility to the unborn child--a feature which is to override the moral consequences of other features of particular cases, e.g., the fact that the pregnancy was the result of rape. What

is not in prominence is the question of what sorts of precautions one would have to take in order to prevent pregnancy due to rape, and whether these precautions would not conflict with other rights one has. Thomson's fantastic analogy brings out just these features, features which, by implication, are not supposed to matter to the issue of responsibility.

But they do matter. Voluntarily opening the window and keeping carpets and upholstered furniture are not privileges, one cannot be expected to refrain from such things. And to show the force of this consequence of the analogy, Thomson reads back through the analogy from the fantastic to the non-fantastic, retranslating and exposing absurdities:

Some may argue that you are responsible for its rooting, that it does have a right to your house, because after all you could have lived out your life with bare floors and furniture, or with sealed windows and doors. But this won't do--for by the same token anyone can avoid a pregnancy due to rape by having a hysterectomy, or anyway by never leaving home without a (reliable!) army. (op. cit.)

But what is the similarity between a case of pregnancy due to rape and a case of rooting due to defective window-screens? Whatever similarity there is here is not a similarity entirely based on property-sharing. Thomson's analogy would not necessarily be improved if the story about people-seeds was made to look more like the way things really are. Her analogy explores conceptual relationships, the role of precaution-taking in the context of the ascription of responsibility, the significance of unwanted, but unavoidable, pregnancy to the question of the moral status of abortion in certain cases. All that is required of Thomson is that she tell a tale that is recognisable, however fantastic. There is similarity here, but it is a similarity of structure, conceptual resemblance not property-sharing.

We can get a grip on the notion of similarity operative here by looking again at the case of the woman who shot her dead husband. Compare that case with one in which someone carefully plans to poison a business rival by putting strychnine in his coffee. But after the victim drinks hardily from his cup, and before the poison takes effect, he swallows a fishbone and strangles to death. This case, I want to claim, is more analogous to the case of the shooting of the dead husband than that case is to the slightly altered version of it where the husband was not dead at the time of the shooting. It is more analogous because we have more reason to call them the same kind of case--attempted murder. The two examples are more analogous because we can get clues about the nature of one by setting it alongside the other. They are more analogous because the consequences of misdescribing the cases as cases of murder are equally absurd, namely that one can, with malice aforethought, unjustifiably kill a human being who is already dead.

What these examples show is, first of all, that the formal logician's pessimistic attitude toward argument from analogy is unjustified. The alternative to the property-sharing analogy is not a muddled and unusable relationship between two things. Indeed, property-sharing is not always a necessary condition for relevant similarity. But what these examples also show is that Cardozo's attitude is equally unwarranted. As Becker has shown, Cardozo himself was relying on analogical reasoning in his decision. Analogical reasoning is reflected in his reason for not being misled by the 'dry logic' of the railroad's argument. For how does he argue? He draws out the consequences of seeing Hynes's behaviour in the light of the behaviour of a trespasser, implicitly suggesting that no equally absurd consequences will follow if we see his conduct in the light

of a user of a public waterway (vide Becker, 253-254). Indeed, the substance of Cardozo's argument is that a property-sharing analogy in this case is unsuitable: the problem of appreciating Hynes's behaviour in the legal context is more like the problem of deciding whether Jones is a Good Samaritan than it is like the problem of deciding whether Brown has what Smith died of.

But can we draw a sharp line between inductive and non-inductive analogical arguments? Earlier on I suggested that judicial reasoning is substantially analogical, and some philosophers have insisted that moral and aesthetic reasoning is also analogical (e.g. Margolis (2)76; Findlay, 250; and Kovesi 114 f.). Is inductive or is non-inductive analogical reasoning the working logic appropriate to these spheres?

One of the points I wish to make here is that there is no sharp line to be drawn between inductive and non-inductive analogical reasoning, and that, as a consequence, one must consider these arguments on a case-by-case basis in the contexts where they do the work they do. There may very well be analogies which depend upon the sharing of explicitly specifiable and precisely statable properties which also function in the dynamic way that the analogies of our last few examples do. An analogy would not exist unless there was similarity between the analogues, and one way, although not the only way, in which two things can be said to be similar is if they share properties. Nonetheless, it has been a theme of this thesis that, for the most part, the sorts of questions which require reflection in the law, in morality, in aesthetics and in philosophy generally, the sorts of questions which involve borderline disputes, which generate moral conflicts, paradox or just perplexity, are questions about concepts. And deductive and inductive reasoning,

being specialised forms of reasoning, presuppose the solutions to such questions. We therefore have reason to suppose that the analogical reasoning which is suitable for the legal, moral, aesthetic and philosophical spheres is not inductive in nature. But the proof of this lies in examples of the questions and examples of reasoning about those questions.

But if we have managed, by considering cases, to begin to uncramp the notion of an analogy, we have also gotten some idea of the issues which are involved in a treatment of reflective reasoning by example. We have still to try to answer them. One of these issues looms large. All three of the attitudes towards argument from analogy mentioned above are pessimistic, and each is pessimistic for the same reason. The value of any analogy is a function of the relevance of the comparison. We cannot even begin to tell the good from the bad analogical argument until we have some way of distinguishing the significant from the insignificant analogy. But relevance and significance are extra-logical considerations, so either (i) we should generally distrust such arguments; (ii) we should make the notion of similarity precise and quantitative and not worry about the counter-intuitive consequences of defining it that way; or (iii) we should leave the entire matter to our (or other people's) intuitions. (Compare: morality is not a determinate system, so either (i') we should be sceptical about every moral assertion; (ii') we should reduce ethical notions to scientific ones and not worry about the counter-intuitive consequences of doing so; or (iii') we should trust our (or other people's) intuitions about moral matters.)

The pessimism here is pessimism over the requirement of legitimacy. And it is justified; it is healthy. As Wisdom notes: "There does at

this point draw on us the shocking affinity in kind between the procedures of respectable reflection and those which we regard as disreputable" ((3)). How are we to distinguish good analogical reasoning, or good reasoning from examples of bad reasoning, reasoning which confuses and proves "black white and sane men crazy" (Ibid.)?

The danger of reasoning from analogy is that we can be misled by superficial resemblance. Moreover, there is, as detractors of analogical reasoning are quick to point out, a lack of logical compulsion in the conclusions of analogical arguments. Even Perelman and Olbrechts-Tyteca, who in their book The New Rhetoric stress the importance and ubiquity of nondemonstrative argumentation, admit that analogical reasoning is an "unstable" form of argument:

For the person who rejects the conclusions will tend to assert that "there is not even an analogy," and will minimize the value of the statement by reducing it to a vague comparison or merely verbal resemblance. (393.)

Jeffrie Murphy in his article "Law Logic" has also argued that whenever analogical argument is used there must be some other, more demonstrative reasoning available to support it; for if the existence of a relevant analogy is denied we cannot, he contends, defend the analogy by means of another analogical argument. To do so would involve us in vicious regress (196). Now this sort of suspicion about argument from analogy, argument from example, and indeed all non-deductive, non-inductive procedures, stems from the prevailing view that despite the genuine understanding and insight that may result, these ways of arguing fall short of what it means to reason.

It is not uncommon, therefore, to find philosophers listing argument from analogy, argument from example and others under the rubric 'Techniques

of Persuasion', or 'Rhetoric'. Deductive reasoning is demonstrative; it is compulsive ('If you accept the truth of the premisses of a valid deductive argument, then you must accept the truth of the conclusion'.). And inductive reasoning, though nondemonstrative, is compulsive in so far as it follows well-defined and intuitively acceptable rules. But argument which is neither deductive nor inductive is not only nondemonstrative and noncompulsive, it is also geared to be a tool of persuasion. Of course, few would want to say that the fact that analogical argument is a technique of persuasion alone counts against it being a form of reasoning. Persuading people is what arguing is about, however one proceeds. What is claimed, however, is this. It is difficult to test the legitimacy of particular analogical arguments--the question of the relevance or significance of the analogy is an extra-logical question. Consequently, it is difficult to distinguish the good arguments from the bad ones, the respectable from the disreputable. The chances of being mislead or confused by analogical argument, therefore, are higher than the chances of being mislead or confused by an invalid deductive argument or an incorrect inductive argument. The risks are higher because the method of arguing is insecure.

Now this argument puts us in a very awkward position: I have been arguing that all reasoning depends for its legitimacy on reflective reasoning, on non-deductive, non-inductive reasoning from case to case. If reflection is the source of legitimacy, and if it is an insecure form of reasoning, or worse yet, not bona fide reasoning at all, then this leads to what Wisdom in the Virginia Lectures from time to time calls the "darkest hour before dawn", namely, the conclusion that nothing can ever be rationally established, that no reasoning is really secure and

dependable. So, before we face the question of how reflective reasoning satisfies the legitimacy requirement, we should look carefully at the risks of rhetoric.

Rhetoric and Argumentation

In the Gorgias, Socrates claims that rhetoric is a form of flattery, a technique of persuasion which is employed to amuse or play with an audience in the hope that they will find a particular side of a question acceptable (463 and 502). To this Callicles, a rhetorician by trade, objects that, although Socrates may have a point given the actions and motivations of the majority of orators who address the Athenian assembly, still rhetoric can surely be used for good as well as ill. Socrates somewhat reluctantly agrees, admitting that rhetoric is of two sorts:

...one, which is mere flattery and disgraceful declamation; the other, which is noble and aims at the training and improvement of the souls of the citizens, and strives to say what is best, whether welcome or unwelcome, to the audience....(503.)

Yet, Socrates insists, we should be wary of the rhetorician. Socrates's worry appears to be about the nature of rhetoric; and here he is drawing on a distinction which is also made in the Theatetus (201) between persuasion which produces opinion and persuasion which produces knowledge. Mindful that all teaching, is so far as it is effective, is persuasion, Socrates claims that only the dialectical technique of teaching counts as knowledge-producing persuasion. But how do rhetoric and dialectics differ?

Two differences are brought out in the Gorgias: (1) Rhetoric is sham dialectics. Socrates argues that rhetoric is to dialectics as cookery is to medicine (464-465). Cooking produces pleasure, but not necessarily health; it is--or can in particular instances be--the

counterfeit of a true art. So too, rhetoric aims primarily to stimulate rather than to teach; it is the counterfeit of an art. (2) Rhetoric is a tool of public speaking. The orator makes long speeches and is constrained by time (he works 'against the clepsydra' in the assembly). The dialectician employs short questions and answers, he is not constrained by time and need not confine himself to what his audience wants, or would find pleasing to hear. This second difference seems to be the crucial one for Socrates. He was greatly impressed by the limitations imposed by the Athenian assembly and was concerned about the distinct lack of Socratic responses, rejoinders and replies in such debates (Gorgias, 229; and Protagoras, 329). The virtue of dialectics is its conversational style. One question is asked at a time and one can easily keep track of the consequence being drawn from previous steps or the analogies which are being offered. This question-and-answer procedure is exemplified in the two Socratic techniques which are frequently abstracted from the Dialogues: refuting an opponent's original thesis (usually in the form of a definition of some sort) by getting him to draw from it a consequence which contradicts it (elenchus); and inferring a universal definition by leading one's opponent through a series of propositions about particular, parallel cases and then generalising (epagoge).

The Socratic elenchus, in its general form, is a deductive procedure: 'Suppose p; then q. But if q then r, then..., then not-p. Therefore not-p'. And most of the epagogic arguments found in the Dialogues are inductive in form--in the Gorgias, for example, Socrates argues that all wants or desires are painful since hunger and thirst (and others he says he could mention) are painful. Nonetheless, Socrates does not always argue

from particular cases to generalisations or definitions in an inductive fashion. Frequently, he sets out, usually painstakingly, a single case, or a series of parallel cases, and reasons analogically to the conclusion that the case he is interested in has some important property. We have already seen this procedure in action when Socrates draws the analogy, rhetoric:dialectics :: cookery:medicine. But all of this is familiar Socratic technique. The point is that Socrates's objections to rhetoric are directed to the deceptive or confusing use of techniques which he employs often enough himself. Surely Socrates himself would agree that although a cook may stimulate the taste buds while starving the body, the conscientious cook will make the body healthy without sacrificing the taste buds.

It is significant that Socrates rarely explicitly identifies rhetoric with eristics or sophistry, the intentional use of invalid arguments, word-play, concealed premisses, and general obfuscation which often results in self-contradictory conclusions, conclusions of the sort which are reported to Socrates in the Euthydemus (e.g. 'The wise are those who learn; but learners are unlearned, and therefore the unlearn learn and not the wise', 276). Socrates is worried about sophistry. And he is worried about long speeches--for good reason: one must always be able to follow the moves of an argument and have the opportunity to object to the premisses which are advanced or assumed. He objects to flowery and--what these days would be called--emotive language. But his objections to rhetoric are not objections to nondemonstrative reasoning as such. He employs analogical reasoning himself. Indeed, the setting out of examples and parallel cases, and the use of imaginary situations and allegories are very important to his dialectical method. His objections are directed at the misuse of these

techniques. There is, then, a three-way distinction in the background: Socrates rails against eristics or sophistry, a false and evil practice which he sometimes, but not always, identifies with rhetoric, and always distinguishes from dialectics.

The same three-way distinction is employed by Aristotle who composed a treatise on each: On Sophistical Refutations, Rhetoric and Topics. In each of these treatises Aristotle is engaged in laying down rules for nondemonstrative reasoning--rules for proper reasoning, for effective reasoning and for identifying and countering improper reasoning. It is important to note, however, that Aristotle's distinction between demonstrative and nondemonstrative reasoning is different from the modern one. The modern distinction--which I have been presupposing--is a distinction between deductive reasoning and non-deductive reasoning, the latter including inductive reasoning as well as analogical reasoning, reasoning from examples, and the rest. For Aristotle the distinction is made, not just in terms of the form of reasoning, but in terms of the epistemological character of the premisses as well. Thus demonstrative reasoning for Aristotle is syllogistic reasoning from premisses which are "true and primary", that is, scientific principles or self-evident truths, premisses which are believed "on the strength not of anything else but of themselves (Topics, 100a 26-b 21). Nondemonstrative reasoning, on the other hand, is first and foremost dialectics.

Dialectics divides into two species: induction and 'reasoning'--a convenient catch-all term which denotes a wide range of procedures, syllogistic, semi-syllogistic (enthymemic syllogisms), as well as procedures for, e.g., discovering ambiguity, determining the degree of difference or sameness between things and properties, and many others.

But, more to the point, dialectical reasoning is reasoning which is based, not on self-evident propositions or scientific-metaphysical first principles, but on opinions which are generally accepted, where "those opinions are 'generally accepted' which are accepted by every one or by the majority or by philosophers--i.e. by all, or the majority, or by the most notable and illustrious of them" (Topics, 100b 20-24). This important aspect of dialectical reasoning is reflected in the title of the treatise where Aristotle discusses it: the topics (topoi) are the place or seats of arguments, the commonplaces or assumptions which form the starting points for discussion--those claims which, for the purposes of the argument, are held to be uncontentious.

Besides dialectical reasoning Aristotle recognises two others--eristic reasoning and 'mis-reasoning'. Eristics differs from dialectics in two ways. The reasoning may start from opinions which, for one reason or another, seem generally accepted or acceptable but are not; or the reasoning may itself be contentious and not make the point it seems to make. In On Sophistical Refutations the distinction is drawn in terms of that quality of arguing which is usually associated with the term 'sophistical', viz. outright deception, reasoning for show: "...for some people it is better worth while to seem to be wise, than to be wise without seeming to be..." (Refutations 165a 19-21 and 171b 33-35). Mis-reasoning is reasoning which, formally speaking, is acceptable enough, although it is based on premisses which are neither true and primary, nor generally accepted, that is, premisses which are both false and misleading (his example is that of an argument in geometry which is based on a misleading diagram). When Aristotle outlines the different ways in which an argument can be eristical, he is specifying what, these

days, are called the formal and material logical fallacies: equivocation, amphiboly, accident, composition and division, many-questions and the rest.

But what of rhetoric? For Aristotle rhetoric is the counterpart of dialectics, it is that art which is concerned with the most effective way of persuading one's audience. Persuasion is, moreover, the orator's demonstration (Rhetoric, 1355a 5-10). Aristotle complains that previous writers on rhetoric have misunderstood their topic; they paid too much attention to forensic rhetoric, that body of techniques which was tailor-made for the Athenian assembly where every citizen was both juror and judge. Oratory techniques concerned what were to Aristotle's mind a relatively unimportant aspect of rhetoric--how to arouse prejudice, pity or anger. What was needed, and what he saw himself as providing, was a systematic account of modes of persuasion. Admittedly, part of that account is an account of the techniques one can use for making one's audience more willing to accept one side of a question rather than the other. But rhetoric is also concerned with proving the true and refuting the false.

To fully appreciate Aristotle's motivation for systematically treating the subject of rhetoric, we need only look at the two reasons he gives for thinking it useful to do so. In the first place, Aristotle (perhaps optimistically) suggests, things that are true and just have a 'natural tendency to prevail' over things false and unjust. It is therefore good, in the long run, to equip people with rhetorical techniques. And secondly, it is often the case that not even the possession of the exactest knowledge will make it easy for some audiences to become convinced of the truth of a true claim or the

falsity of a false claim (Rhetoric, 1355a 20-26). If it is difficult to accept the first of these reasons why knowledge of modes of persuasion is useful, then, at any rate, the second should suffice. Rhetoric is useful because it is a commonplace that the true and the just are not always convincing while the false and the unjust are not always unconvincing.

Aristotle is in his treatment of rhetoric trying to allay the fears which Socrates voiced. For one surely must distinguish between the possession of the rhetorical talent and the improper, because unjust, use of it. At one point, Aristotle puts the question this way:

What makes a man a 'sophist' is not his faculty, but his moral purpose. In rhetoric, however, the term 'rhetorician' may describe either the speaker's knowledge of the art, or his moral purpose. In dialectic it is different: a man is a 'sophist' because he has a certain kind of moral purpose, a 'dialectician' in respect not of his moral purpose, but of his faculty. (Rhetoric, 1355b 17-22.)

Aristotle is making two important points here. (And he is making them in an important way, by mixing talk about rational faculties with talk about moral purposes. I shall be relying on and developing this mixed mode of expression in the next section.) The first is that to call a person a rhetorician, or an argument an example of rhetoric, is not necessarily to impugn the rationality of the person or the argument. Other considerations are relevant to the question whether a rhetorician is a sophist or a piece of rhetoric sophistry. These considerations have in the main to do with the intention of the person who puts forth the case--is his rhetoric designed to deceive or to instruct? The purpose of rhetoric as such is not simply to succeed in persuading, it is rather "to discover the means of coming as near such success as the circumstances of each particular case allow" (Ibid., 1355b 9-12). The purpose of rhetoric, in short, is to present as clearly and effectively

as one can one side of a question. But this is advocacy, putting forward the best possible case on both sides of a question so that both sides get a fair hearing.

The second point Aristotle is making here takes him some distance from the Socratic position with regard to the built-in virtue of dialectics: Aristotle is acknowledging that to call someone a dialectician is not thereby to preclude the possibility that his arguments are, and are intended to be, deceptive (cf. Ibid., 1355b 3-8). Someone is a dialectician in respect of his faculty of reasoning in a nondemonstrative fashion towards, generally speaking, the refutation of a claim. If the dialectician has a firm grasp of the techniques of persuasion as well, then he is a rhetorician, and as such a more effective dialectician. But his talents may be employed to deceive. This leads to the wider point that any form of reasoning can be employed to deceive and confuse: the sophist can equally well lie with statistics as with argumentum ad hominem.

If we follow Aristotle's lead on the nature of rhetoric, then the moral we should draw with regard to the worry expressed at the end of the last section is this: The question of the legitimacy of non-deductive, non-inductive reflective reasoning is not answered by saying that such reasoning is, or is like, the use of rhetorical techniques of persuasion. We cannot, that is to say, decide which instances of analogical or case-by-case reasoning are instances of good or of bad reasoning simply by applying the term 'rhetorical' to some while withholding it from others. But we should, nonetheless, recall the reasons why Socrates in the Gorgias opposes rhetoric with such vigour. One should be suspicious of, if not completely reject, a technique designed to be used by anyone who is impatient to arrive at the truth and wishes instead to arrive rapidly

at conclusions which will, by the mere mode of their presentation, convince and impress. With this sort of motivation, the chances are indeed slim that the resulting arguments will be rationally credible.

It should be recalled further, however, that in the Phaedrus Socrates does mention the possibility of a rhetoric suitable for the philosopher, a rhetoric which would be capable of convincing the gods themselves (273-274). In the sense in which rhetoric would be capable of this sort of credibility, it would be rational. It would be, indeed, the sort of rhetoric which Aristotle has in mind, a rhetoric which is not opposed to dialectics but a counterpart of it. The difference between a sophistical refutation and a dialectical refutation, Aristotle tells us, is that whereas the former is refutation aimed at and relative to some one person, the latter is absolute, it is refutation for anyone (Refutations, 170a 12-13). Likewise, the difference between sophistical persuasion and rhetorical persuasion is the difference between causing someone to change his mind or come to an opinion and presenting arguments which are both effective in persuading and worthy of persuading anyone.

It may be impossible to release the term 'rhetoric' from the hold which current usage has over it, usage which emphasises the elements of insincerity, of excessive, exaggerative and abusive language. But one recent work has tried to return to Aristotle's conception of rhetoric by emphasising the direct link which he saw between it and dialectics. I am thinking here of Perelman and Olbrechts-Tyteca's The New Rhetoric: A Treatise on Argumentation. The close similarity between these authors' approach to the nature of reasoning and that followed in this thesis might best be brought out by the following question which they ask at

the outset: Given the evolution of logic toward formal deductive calculi, and granting the real advances that have been made there, still

must we draw...the conclusion that reason is entirely incompetent in those areas which elude calculation and that, where neither experiment nor logical deduction is in a position to furnish the solution of a problem, we can but abandon ourselves to irrational forces, instincts, suggestion, or even violence? (3.)

Their answer is that reason is not incompetent or ineffectual in such areas and for such problems. Their answer takes the form of a theory of argumentation, a study of the "discursive techniques allowing us to induce or to increase the mind's adherence to the theses presented for its assent" (4).

In several places, Perelman argues that modern logicians have whole-heartedly and uncritically adopted a conception of proof, and a correlative conception of knowledge, which necessitates that genuine reasoning must be demonstrative (in the modern sense, not Aristotle's). The culprit, although it might be unfair to place all the blame on one philosopher, is Descartes; the source of this error is Cartesian epistemology (vide Perelman's "Self-Evidence and Proof" in Perelman (1); and Perelman (2)). Perelman points in particular to Descartes's motto in The Discourse on Method, 'take well nigh for false everything which is only plausible'. When applied to the general problem of knowledge, the result is that one must discover a self-evident basis ('clear and distinct ideas', sense-data, axioms) upon which the rest of what we can truly say we know can be built. Reasoning on this view is secure when it allows one to move necessarily from self-evident propositions to other propositions; genuine reasoning is compulsive and conclusive. The Cartesian programme produced the invidious dichotomy which set as the only alternative to absolute certainty, absolute scepticism.

The Cartesian programme, Perelman argues, effectively relegates argumentation and deliberation to the sphere of the non- or irrational. For it is the very nature of argumentation and deliberation that they deal not with the certain but with the plausible, the credible and the probable (cf. Aristotle, Nicomachean Ethics 1112a 30-1139a 13). Indeed, the very notion of a dispute is altered on this view to mean, not a conflict between people who may be proceeding rationally in their attempt to establish one or the other side of a question, but a conflict between reason and non-reason; for on the Cartesian view

if two men defend opposing theses on the same issue, at least one of them is irrational, since he is necessarily mistaken; all disagreement is a sign of error and proves that seriousness is lacking. (Perelman (1), 17.)

The difficulty with this view is not merely that genuine disagreement becomes impossible, but also that, as we have seen before, no rational support for the 'self-evident' basis of knowledge can be forthcoming. What was intended to be the securest possible structure of knowledge becomes instead the easiest target for the sceptic. On this view, no reason, no argument is possible in favour of the foundation of the entire edifice.

The rejection of the Cartesian picture of reasoning and knowledge--the exemplification of Deductivism in epistemology--requires the rejection of the claim that there are propositions whose truth is self-evident, propositions for which no justification is necessary, and, more importantly, for which no justification is possible. This is indeed Perelman's strategy. But he is alive to the fact that one must do more than simply deny that the only justification there is is deductive in form, the only proof deductive proof; and one must do more than simply deny that there are such

things as self-evident truths. One must also offer an alternative account. That account is his theory of argumentation, his 'new rhetoric'.

One of the keystones of this theory is the thesis that the very possibility of human knowledge, and of an intellectual life, requires

...the assumption that there are reasons which can be found for undertaking to believe something without these reasons being so dazzlingly self-evident that the propositions believed in stamp themselves as truths on the mind of any rational being who attentively studies them. (Perelman (1), 119.)

If one holds to the Cartesian dichotomy which entails that the only alternative to absolute certainty is absolute doubt, then one is forced to oscillate between dogmatism and doubt. For one is committed to the self-evidence of a set of propositions while being, at the same time, unable to argue for them. What is needed is both a wider conception of what can count as a basis for reasoning and a wider conception of reasoning. But both of these conceptions are explicit in Aristotle's account of dialectics. As we saw, dialectical reasoning is based, not on self-evident propositions, but on other kinds of topoi, common-places or opinions that are accepted 'by all, or the majority, or by the most notable and illustrious of them'. And dialectical reasoning is not, in Aristotle's sense, 'analytical' or apodeictic; it deals with reasons which convince without being completely compelling, it does not yield truths but probabilities. One can, without self-contradiction, assent to the premisses of a dialectical argument while refusing to assent to the conclusion. Dialectical reasoning provides reasons which, though separately inconclusive, are nonetheless severally cooperative in favour of a conclusion; reasons which, as Wisdom suggests, "are like the legs of a chair, not the links of a chain" (Wisdom (2), 157). Moreover, as Aristotle saw, the existence of dialectical reasoning provides us

with a means for rationally underwriting the self-evident truths or scientific first principles which form the basis for demonstrative reasoning. For although principles of these sorts are the prius of all scientific knowledge,

...it is through the opinions generally held on the particular points that these [principles] have to be discussed, and this task belongs properly, or most appropriately, to dialectic: for dialectic is a process of criticism wherein lies the path to the principles of all inquiries. (Topics, 101b 1-4.)

In The New Rhetoric, Perelman and Olbrechts-Tyteca develop in massive detail the Aristotelian conception of dialectics from the perspective of a general theory of argumentation. In the process, as I have already noted, they hope to re-establish the thesis, so central to Aristotle's picture of the nature of nondemonstrative reasoning, that a theory of the discursive means of obtaining the audience's adherence to claims offered up for assent and dialectics are closely connected subjects. Their procedure, in rough outline, is to consider the structure of rhetorical argument first with reference to the various starting points of argumentation--the topoi, considered as general areas of agreement as well as particular premisses--and second with reference to the various techniques or schemes of argumentation. Both of these approaches are worth exploring in detail.

Argumentation is fundamentally topoi-orientated--that is to say, not merely axiom-orientated--because it is always directed to an audience and is always designed to persuade. In order to persuade an audience--and the authors provide for the case where the audience and the speaker are one, i.e., the case of deliberation--one must begin from a basis of agreement. From that starting point, the object is to carry the audience along step by step on the basis of new levels of agreement until, finally, agreement is reached over the thesis which is being argued. In order for

argument to begin, however, some common ground must be found. But that common ground need not be--as in the deductive model--some set of axioms, nor need it be--as in the inductive model--primary, uncontentious data. All that is required for a starting point is something which, as Perelman puts it, "can provisionally be accepted without discussion" ((1), 119). But what is provisionally accepted for one argument may very well be the very heart of the controversy in another argument. In order to argue about a matter which is in dispute one must begin with that which is not in dispute. But one can set down what is not in dispute without thereby crediting it with self-evidence or logical/epistemological priority.

The following analogy suggests itself: If a stranger asks directions to the Post Office our first task is to find some place the stranger does know how to find. We need a starting point. If he knows how to get to the Courthouse we might be able to help him out simply by informing him that the Post Office is on the same street, just south one block. But if the Courthouse is the appropriate starting point in this case, it may not be in another: it is not the place where everyone must start in order to find their way around. The next stranger may want to find the Courthouse and know where the Post Office is. The same story can be told in the case of argumentation, where our problem is finding our way around logical space.

There are, of course, various sorts of starting points. The authors distinguish between topoi which take the form of premisses specifying a single starting place for argument and those which designate starting points of relative generality. These latter topoi, in the Aristotelian tradition, serve as headings under which arguments can be grouped: topoi

of quantity, quality, order, the existing, essence and the person (The New Rhetoric, 83-99). For example, an argument founded on, or guided by the view that 'a greater number of good things is more desirable than a smaller' (cf. Topics, 117a 16), is an argument which is founded on and guided by the topos of quantity. There are countless, specific topoi of quantity as there are countless ways in which the desirability, superiority or worth of two things can be granted to be open to quantitative consideration. Moreover, as the authors readily admit, it would be a very difficult task indeed to provide an exhaustive list of all the topoi which could be used (Ibid., 85). Indeed, one of the primary differences between the 'new rhetoric' and the Aristotelian dialectics upon which it is based is that whereas the latter was intended to be a complete study, involving an exhaustive classification of both topoi and argument schemes, the former is limited to a survey of the clear and ready-to-hand examples of topoi and schemes, a survey which gives the reader an idea of the indefinite number and variety of each (Ibid. 509 ff.).

Of considerably more interest to our present concern are the specific topoi. Here the authors identify as initial objects of agreement facts, truths, presumptions, values and hierarchies (67-83). By 'facts' and 'truths' the authors intend whatever claim--empirical, a priori, scientific, religious or philosophical--which, for the purpose of an argument, is not contested (but could, in another argument, be the object of controversy). Likewise, presumptions (of, e.g., innocence, the reliability of witnesses, &c.), values and hierarchies (e.g. superiority of men over animals) are claims which for the purposes of particular arguments are not in dispute. Specific topoi are thus initial premisses of an argument, explicitly stated or assumed without comment.

It is easy to see how specific topoi figure in analogical arguments, and how, for that matter, the existence of a starting point of some sort is essential to the construction of an analogy. Consider Aristotle's figurative analogy, "For as the eyes of bats are to the blaze of day, so is the reason in our soul to the things which are by nature most evident of all". Here the commonplace which is being presupposed and exploited is that of the phenomenon of bats being blinded by, and hence being unable to clearly make out, the world brightly lit up. If Aristotle had said instead, 'For as the Ultimate ponders the perceptual penumbra, so is the reason...', we would have a difficult time trying to see the point of the analogy. We would lack a commonplace, some relationship which we do understand alongside which the relationship which we do not yet understand is being held. Clearly then, analogical reasoning requires a topos of some sort.

The notion of an argumentative topos can also be employed to add support to the alternative jurisprudential view which we introduced in Chapter III. It will be recalled that both Formalism and Realism had difficulty in accounting for the function of precedent cases in the common law. The difficulty was this. How could the ratio decidendi of a case be binding on future cases if (i) the ratio cannot satisfy the logical role of a principle from which decisions about future cases can be deduced, and (ii) if there is no single, unambiguous ratio of a case? But, on the other hand, if the Realist is correct in saying that rationes have no effect on future cases, and that stare decisis is a principle without application in the common law, then how can we account for the fact of unity and continuity there?

It has been argued by Julius Stone that the ratio decidendi of a

case serves the function of a starting place for future argument concerning new cases. Moreover,

...it may be that when we call a case "a leading case", what we really mean is that that case is, for the time being, a seat of argument. What precisely the case stands for can rarely...be determined by analysis of either the facts or the judgements. Yet some composite of these is still where argument tends to start, once it becomes common ground that the case is a leading case, and as long as it remains so. In that sense a leading case is a topos of legal argument. (335.)

If we think of a precedent, or a group of competing precedents, as starting points for lawyers' arguments and judicial reasoning, then we can orient the judicial process. Stone spends several chapters of his Legal System and Lawyers' Reasonings showing by a consideration of cases that the Formalist's 'first principles of law' represent in fact categories of "illusory reference", categories which appear to lead judges by compelling and stringent deduction to their decisions but which in fact leave them free to choose between different legally justifiable results. There are, built into the system of precedent, "leeways of choice" which arise from (i) the nondeterminate nature of legal and legally relevant concepts used in statutes; (ii) the various competing versions of the ratio decidendi of previous cases; and (iii) the various competing judgements which have grown out of a single leading case (Ibid., 274-278).

We have already seen how these and other factors taken together argue against the possibility of an adequate, purely deductive characterisation of judicial reasoning. But if we characterise judicial reasoning as dialectical--or rhetorical in the sense in which Aristotle and Perelman use that term--we shall, Stone believes, more clearly understand the actual process of the common law. As against both the Formalist and the Realist, we can see that by insisting on the dialectical model

of judicial reasoning, "even logically uncompelled choices [can be] made with reasons publicly examinable" (Ibid., 326). In particular, the notion of a topos of an argument changes the complexion of judicial reasoning: instead of the data from which judges begin--the premisses of legal arguments--being legal propositions able to serve the logical role of major premisses in deductive arguments, this data can be seen as a composite of such determinate concepts as there are in the law with considerations of justice or policy. This is the judicial starting point.

Yet, despite the intrinsic philosophical importance of the notion of a topos, Perelman and Olbrechts-Tyteca are more concerned in their treatise to lay out systematically, although, once again, not exhaustively, the various schemes of argumentation. Discussion of these techniques of reasoning constitute the bulk of their work, and it would be pointless to attempt to treat separately the hundred-odd distinct techniques they consider. In large part, the cases they consider and the examples they give speak for themselves and I could add little by rehearsing them. What I would like to do, however, is to treat two general points, two classificatory rationales, which come out of this work. These points bear directly on issues which I have considered only briefly and schematically above.

First of all, the authors take care to separate off a set of argumentative techniques which bear a resemblance to deductive modes of reasoning. These rhetorical arguments, the authors contend, "lay claim to a certain power of conviction, in the degree that they claim to be similar to the formal reasoning of logic or mathematics" (The New Rhetoric, 193). Given this similarity, the authors call these rhetorical arguments 'quasi-logical'. But they are quick to add that they do not

intend by using this name to imply that formal, deductive reasoning is primary and argumentation an approximation, or an imperfect form of it. Rather, "formal reasoning results from a process of simplification which is possible only under special conditions, within isolated and limited systems" (loc. cit.). This is, of course, one of the things which I have been trying to show in the preceding chapters, although I have chosen to put the point in another form: formal reasoning is specialised reasoning, it is reasoning the undeniable power and perspicuity of which is the result of a particularly favourable conceptual environment. The authors are interested in reasoning which takes place outside of closed, determinate systems, as I have been. And their interest in quasi-logical arguments is not motivated by a concern to see how arguments can fail to be rational--by falling short of the standards of deductive reasoning--but rather by a concern to see how arguments in law, morality, philosophy and elsewhere can be rational without being either deductive or inductive in form.

It is in this spirit that the authors consider the quasi-logical argument patterns which involve incompatibility. With deductive reasoning, to build an argument based on a contradiction--say a reductio ad absurdum--requires that one show two or more propositions to be explicitly inconsistent. But in ordinary contexts, where the predicates employed are not always univocal, where borderline cases are acknowledged, we rarely can isolate explicit contradictions. Indeed, as the authors point out, in normal argumentative contexts, when someone asserts a proposition and its negation simultaneously, we are inclined to look for some way of interpreting what he says in order to avoid the obvious, but uninteresting, inconsistency (Ibid., 195).

Suppose, for example, someone were to claim in a dispute over Hitler's status as a moral being, 'Certainly, he was a human being; but at the same time he wasn't'. Now if someone were to reply, 'But that's absurd, no one can both be and not be a human being', it is likely that he would be open to the charge of playing the pedant. The obvious reply would be rather, 'What do you mean; what distinction are you trying to make?'. And that question can be answered in a rational and philosophically interesting way. We need only recall that the abortion debate rests in part on the question of whether something which counts as a human being considered from one perspective (the offspring of two human beings) can also be considered as something which does not count as a human being (part of the mother's body).

Incompatibility results when it becomes obvious that a choice must be made between one way of talking or acting and another. Incompatibility is a sign of a conflict which requires resolution. The conflict may involve the application of two different precedents to a single instant case, two competing theories about a particular phenomenon, or the application of a single concept to two different cases. Indeed, the most natural way of characterising moral and legal disputes is in terms of the incompatibility of two or more assertions.

And conflicts stand in need of argumentation. There are various ways in which one can proceed in order to attempt to resolve incompatibility. One can draw out of the claim one's opponent is making other claims which he would also have to make in the hope of coming upon a claim which he refused to assent to. One can, of course, do this deductively--employing the Socratic elenchus mentioned above. But one need not proceed deductively. One can proceed analogically taking the

claim one's opponent wishes to make about a certain thing or situation and testing it out against similar things or situations, determining in the process what one's opponent is committed to by his original claim.

The second general point that comes out of the treatment of schemes of argumentation in The New Rhetoric is this. The authors maintain that each of these schemes can be characterised either as a process of association or as a process of dissociation. Processes of association are those "which bring separate elements together and allow us to establish a unity among them, which aims either at organizing them or at evaluating them, positively or negatively, by means of one another" (190). Processes of dissociation, on the other hand, are "techniques of separation which have the purpose of dissociating, separating, disuniting elements which are regarded as forming a whole or at least a unified group within some system of thought" (loc. cit.). Schemes of argumentation, in other words, are techniques which expose or bring into prominence similarities or differences, techniques of comparison or contrast. Argumentation itself is seen as an exchange where, on the basis of some topos, competing views on a particular question are argued for by exposing or bringing into prominence similarities or differences between what is at issue, and what, for the purposes of the argument, is not at issue. One attempts to counteract or nullify a proposed unity by trying to expose a more significant disunity; and one attempts to counter a proposed disunity by bringing to light a more important unity.

The primary use the authors make of the association/dissociation distinction is classificatory. But these two notions are conceptually linked to what Perelman elsewhere calls the 'rule of justice', namely "things in the same essential category should be treated in the same way"

(Perelman (1), 15). Thus, if it can be shown that two separate elements can be associated to establish a significant unity, then that counts as a reason for treating these elements in the same way. Conversely, establishing significant disunity counts as a reason for treating these elements in different ways. What these elements are depends on the nature of the object of the dispute; but, typically, we are interested in a question which has to do with how some thing, action, situation, characteristic or feature of some thing, action or situation should be described, i.e. what sort of thing, action, situation it is. The rule of justice requires us to treat like cases alike, different cases differently. Hence the aim of associating and dissociating.

It should be emphasised, however, that for Perelman this rule of justice is also the principal rule of reason, it is, indeed, what it means to be rational (The New Rhetoric, 218-220; and Perelman (1), 79-87). The interpenetration between the normative aspect of rationality and the rational aspect of justice is a dominant theme of Perelman's work. In "The Rule of Justice", for example, he writes that

...if there exists any practical employment of reason, it must manifest itself in just action bearing witness to a rationality which unjust behaviour would not display. (Perelman (1), 79.)

One is reminded here of Aristotle's treatment, in the Nicomachean Ethics, of practical wisdom and moral virtue. For Aristotle practical wisdom is "the quality of mind concerned with things just and noble and good for man" (1143b 21-22). Practical wisdom is a mental talent for deliberating well about what is good, an intellectual virtue which is demonstrated in knowing what sorts of things are conducive to the good life in general (1140a 24-1140b 30 and 1143b 21-28). But for Aristotle, having the ability to deliberate well over matters of human good means that one has

as well the ability to do good things, since "states of character arise out of like activities" (1103b 21-22). The connection between being practically wise, being rational, and doing what is in accordance with practical wisdom, being just, is perhaps best brought out when Aristotle claims: "It is clear, then,...that it is not possible to be good in the strict sense without practical wisdom, nor practically wise without moral virtue" (1144b 30-31). There is, therefore, a strong Aristotelian flavour to Perelman's notion of the rule of justice.

The rule of justice and the notion of an argumentative topos are also closely linked for Perelman. Part of the reason for this is that Perelman conceives of the rule of justice on the model of stare decisis (vide, e.g., Perelman (1), 63). That is, the rule of justice is interpreted to mean that it is rational and just that decisions taken in previous cases should likewise be taken in similar, present and future cases. As we have already noted, it is quite natural to think of precedents in the legal context as topoi, for they are starting places--possibly competing starting places--for reasoning about instant cases. In line with this interpretation of the rule 'Treat like cases alike', Perelman emphasises the importance of a tradition of previous cases, a tradition the justification for which can be taken for granted. One stands by previous decisions until one is given a reason for abandoning or altering them:

It is tradition that is taken as a starting-place, it is tradition that is criticized and it is tradition that is maintained in so far as no reason is seen for departing from it. And this holds good in the most diverse fields--ethics or law, science or philosophy. (Perelman (1), 86.)

Here too there is an Aristotelian flavour to Perelman's remarks, for Aristotle also viewed the topoi as traditional; they are opinions which

are 'generally accepted', accepted by all, the majority, or by the most notable and illustrious. But this emphasis on tradition need not be taken as a stand in favour of legal, moral, scientific or philosophical conservatism. For the point is simply that one begins with what is not controversial, with what is familiar. For as it happens, often the best starting place is just where you happen to find yourself.

Having now explored in some detail the theoretical background for the new rhetoric, we are in a better position to consider the question of the legitimacy of non-deductive, non-inductive reflective reasoning. We have a couple of clues to work with. In the first place, it should be clear that the mere fact that an argument is dialectical--or in the Aristotlean sense, rhetorical--does not mean that it is illegitimate, that it is a bad or a disreputable argument. This remark, probably a commonplace for Aristotle, today requires argument. Perelman and Olbrechts-Tyteca have, it seems to me, supplied that argument. Their work shows that in order to discriminate between good and bad rhetorical arguments an understanding of the structure of argumentation, of non-deductive and non-inductive reasoning, is required.

Secondly, Perelman and Olbrechts-Tyteca's closer look into the structure of argumentation reveals two distinct aspects, and suggests two distinct approaches to the question of legitimacy. We have both to consider the starting places, the topoi of arguments and the schemes or patterns of arguments. The authors are surely correct to emphasis these two aspects of argumentation, for the most conspicuous difference between a formal inference and, in the wide sense, an argument is that whereas in the former it is the logical link which is central, in the latter it is both the link and the character of the premisses, the starting points,

which are important and open to scrutiny. Consequently, the question of legitimacy appears to break into two parts. In considering whether some particular argument is respectable or disreputable we need to consider both the starting place and the scheme of reasoning.

Now, if the standard of legitimacy in the case of argumentation was either deductive validity or inductive correctness, every piece of non-deductive, non-inductive reasoning--which is to say every example of argumentation Perelman and Olbrechts-Tyteca consider--would be illegitimate. But, as our earlier discussions have shown, if validity and inductive correctness were our sole bases for reasonableness, we would be forced to the conclusion that we ought to be sceptical about the very possibility of there being any sort of legitimate and fruitful reasoning. For in the end, deductive and inductive reasoning presuppose the legitimacy of the sort of reasoning which the authors of The New Rhetoric are trying to classify and understand--that is, non-deductive, non-inductive, reflective reasoning. In our search for the basis of legitimacy of reflective reasoning, therefore, we must look beyond the standards of deductive validity and inductive correctness.

The question now becomes, What should take the place of validity and correctness? H. L. A. Hart, in his introductory remarks to Perelman's The Idea of Justice and the Problem of Argument, has written

once we command a clear and detailed view of the way in which in living discourse conclusions are supported by arguments and arguments are evaluated as 'strong' or 'weak', or as having or lacking 'force' or 'weight', it seems dogmatic to confine the term reasoning to logical demonstration or inductive generalisation or the apprehension of self-evident truths. (x.)

I am very sympathetic to the point Hart is making here; but it is not without its problems. Are we to say that there is a standard of legitimacy

for non-deductive, non-inductive reasoning, and that that standard is strength, or force, or weight? We are once again facing Wisdom's 'darkest hour before the dawn'. For if legitimacy is a question both of the nature of the starting point and of the character of the mode of reasoning, and if we disabuse ourselves of the illusion that there are self-evident starting places, on the one hand, and admit that we lack formal characterisations of patterns of reasoning on the other, then are we not opening the door to the sceptic? How are we to evaluate arguments in terms of their strength or weakness, their having or lacking force or weight. We want to insist that every topoi may itself be the object of controversy, and that every argumentative scheme can be exemplified by respectable as well as disreputable arguments. But then, even if it is granted that reflective reasoning is more fundamental than, and is presupposed by, deductive and inductive reasoning, are we nonetheless left adrift on the question of what constitutes the legitimacy of particular examples of reflective reasoning?

We seem once again to be staring at the three pessimistic attitudes toward nondemonstrative reasoning which were brought out in the previous section: Are we not forced to admit either that we must avoid such reasoning, or that we should change it to fit into the moulds of deduction or induction, or that we should leave the whole question of when arguments are legitimate to our intuitions? I would like now to show how the notion of legitimacy for reflective arguments can be developed so that each of these attitudes can be set aside in favour of a more optimistic one.

Reflection and Legitimacy

At one point in Proof and Explanation, Wisdom considers the question

of how one would distinguish a good case-by-case procedure from a bad one. His immediate answer is, "By a case-by-case procedure". Later he remarks that this question is crucially important since it leads to the first step on the road to the sceptic's thesis that there is no good reasoning, no possibility of proof at all. For the question can be put this way: Given that all reasoning comes in the end to case-by-case reasoning, "...how do we discriminate between good and bad reasoning, thinking, application of the understanding?" Well, Wisdom answers, not by appealing to a list of rules which give the necessary and sufficient conditions for argumentative goodness. There is no such list. To be sure, there are rules to which we can appeal when we consider the special question of whether some deductive argument is valid. But even here the rules may not do the job, for we can still question the rules themselves. But if we do question the rules then, in the end, we must resort to examples. But if there are no rules, then there is no criterion of legitimacy for case-by-case reasoning. And if there is no criterion of legitimacy, how will we be able to separate the legitimate from the illegitimate, or more importantly (given the spectre of scepticism), how will we be sure that there are any legitimate examples of case-by-case reasoning? Wisdom's answer: we need to look at instances of reasoning which are good, instances which are bad, and instances where the reasoning is in some respects good and in other respects bad.

It is easy enough to formulate the sort of response which would spring to the lips of most philosophers. At best, Wisdom is guilty here of begging the question. If reasoning by the presentation of examples were a legitimate procedure of reasoning, then arguing that the question of the legitimacy of case-by-case reasoning rests on a consideration of

cases would be legitimate, although otiose. But, the issue is whether any instance of reasoning by example is legitimate. So at best Wisdom is reasoning circularly, using a procedure in order to show the legitimacy of that procedure. But at worst, Wisdom is not even addressing himself to the question. For what reason do we have for thinking that a consideration of examples of good and bad reasoning from examples will provide us with a criterion of legitimacy, or if not a criterion, then at least some reason for thinking that there are cases of good, of legitimate reasoning from examples? Unless Wisdom has up his sleeve some sort of intuitionistic view, we have no reason to think that his examples have anything at all to do with the issue of legitimacy.

There is no denying that on the face of it these are strong objections to what Wisdom is about in Proof and Explanation and elsewhere. Nonetheless, it seems to me that he is correct: to distinguish the legitimate from the illegitimate we must proceed by means of a consideration of cases. What I want to do, however, is to show why Wisdom's answer is correct. His answer is at the heart of the entire project of characterising fundamental reasoning as non-deductive, non-inductive reflective reasoning. But his answer needs some filling out; a wider story must be told.

To get at that wider story I want to consider some other ways of looking at the question of the legitimacy of nondemonstrative and non-inductive reasoning. In effect, I want to look at some ways of understanding the legitimacy of reflective reasoning which are bad, and some other ways which are in some respects bad and in other respects good.

Suppose we begin with the worst of the lot. It might be suggested that the only way in which we can understand the legitimacy of nondemonstrative reasoning is as an approximation to deductive validity. We

have already encountered this suggestion when a couple of sections back we considered the formal logician's favoured conception of argument from analogy. In terms of analogical reasoning this 'approximate-validity' approach goes like this. We first grade analogies on a scale given by the 'shared-property' conception of analogy, such that two pairs of analogues can be compared for strength or closeness simply by means of counting up the number of properties shared by the two objects of comparison in each pair. Arguments employing analogies can then be graded, the best possible analogical argument being that argument which employs an analogy between 'two' things which share all their properties, i.e., are identical. The best possible analogical argument on this view is, of course, a deductively valid argument whose major premiss is an identity claim. The legitimacy of a particular analogical argument is thus conceived to be a matter of degree, optimal legitimacy being deductive validity.

I have already spent some time arguing against the property-sharing conception of analogy, concluding that that conception has only a limited application and is in any event derivative from a notion of analogy which does not reduce to property-sharing. We have, therefore, seen the mischief which forcing reflective reasoning into a deductive mould can cause. It is important to see, however, how far one must go in order to construe nondemonstrative legitimacy as approximate deductive validity.

To illustrate the problems with this approach, I want to talk briefly about some features of moral argumentation. First, a distinction: Among moral assertions there are--besides particular assertions about particular people, acts, places and times--general and (what I want to call) indefinite moral assertions. General moral assertions--those

usually called moral principles or moral axioms--are typically assumed to be universally quantified statements, e.g., 'Murder is morally wrong', or universal, categorical commands, 'Thou shalt not steal'. Indefinite moral assertions, on the other hand, are assertions which are indefinitely quantified, assertions which are expressed in what might be called 'for the most part' quantificational form (cf. Ross's claims about the use in moral reasoning of "intuitive induction" to moral rules of thumb, or media axiomata, (2) 168-174). Indefinite moral assertions--'Those who refuse to hire Blacks solely because of their race tend to perpetuate injustice', or 'A person who willingly assists in an abortion can be thought to be an accessory to murder'--sum up in a perspicuous manner particular cases. Their use in moral reasoning is primarily that of redirecting our attention, forcing us to look at a particular act or a particular kind of act in a certain way (as unjust, or as an act which assists someone in committing a murder). Unlike universally quantified, general moral assertions, an indefinite moral assertion is about some indefinite collection of particular cases all of which warrant the same sort of judgement. But because of the logical form of these indefinite assertions, they cannot serve the role of major premisses in deductive arguments, the role general moral assertions are designed to serve. It is, moreover, in the spirit of Deductivism to abandon such indefinite assertions in favour of the deductively more workable universally quantified assertions.

But this puts the Deductivist in a familiar dilemma, a dilemma which Wisdom hints at when he says "the only ethical principles worth attention are ones which are false..." (Wisdom (1), 223). That is, moral principles of great generality, say 'Murder is morally wrong'

(or even the sort of principle which Wittgenstein would have us call a grammatical remark--'One ought not to do what is morally wrong'), although they admit of few, or no exceptions are of limited argumentative value. Indeed, it is very difficult to imagine a place for them in moral reasoning. Typically, the object of controversy in a moral dispute is located at a level which is quite different from that at which these highly general principles can be found. The problem is that of describing the act, determining whether, for example, it is or is not correct to call it an act of murder. Moral disputes are usually directed to particular cases, the object of controversy concerns what ought or ought not to be done in a particular set of circumstances. But the claim that all murders are morally wrong does little or nothing to help moral disputants with what is at issue in particular cases. (Very general moral assertions may help to educate someone, to tell a child what it means to say that an act is morally wrong, for example. But outside of an educative context, such a moral principle is of little value.) On the other hand, general moral assertions which are worth attention are usually those which admit of many exceptions.

Consider, for example, the claim that every abortion is morally wrong, a claim which is interesting in an argumentative context precisely because it seems so obviously false. We can easily imagine cases where abortion would be, in the circumstances, morally correct, or at least morally permissible (for example, a case where abortion would be the only means of saving the life of the mother). The false moral principles are interesting because by seeing that they are false we begin to develop a better understanding of the moral nature of the act in question. We begin to see the kinds of overriding considerations there are. But the

role which a false moral principle plays--primarily that of drawing our attention to counter-examples--is not the role which a major premiss in a deductive argument plays. So, in either case, the deductive inference which the Deductivist wants us to focus on, the inference he views as central to moral argumentation, drops out of the picture.

But if we drop the prejudice in favour of the deductive model of arguing from universally quantified principles, the role of moral principles in moral argumentation becomes clearer. Granted that general moral principles can pack a considerable argumentative punch, they are nonetheless always either extremely vulnerable to counter-example or else analytic and of little argumentative value in the long run. In the context of a moral dispute, I want to suggest, when we make assertions which are generalisations and are not interested in the shock-value of a bald, universal claim, we usually intend to make indefinite assertions. We sacrifice the blatant 'all' or 'every' claim for the tempered 'most' claim. When, for example, we claim that abortion is morally permissible, what we usually intend to assert is that the reasons which can be given for an abortion in a particular case tend to outweigh the reasons which can be given against an abortion. And that an indefinite moral assertion can count as a reason on one side or the other of a moral issue is a consequence of the fact that such an assertion is a way of bringing to the argument an indefinite number of particular examples. These particular examples, described in detail, count as reasons. The indefinite assertion is a shorthand way of bringing these reasons to bear on the moral issue of abortion. Similarly, to say that abortion is prima facie morally permissible is just to make the indefinite assertion that there are particular examples waiting in the wings, coupled with the rider that

in light of these examples, the burden of proof can more reasonably be seen to rest on one who would deny that a particular case of abortion is morally permissible.

Once the deductive pattern of reasoning is abandoned the problem of 'reaching higher' for more general principles to justify the principles employed as major premisses vanishes. In the case of moral argumentation, the presumption need never be made that every generalised assertion is a claim about all cases of a certain sort. The indefinite moral assertion, the rational man's moral principle, is granted from the start to be open to counter-example. But a single counter-example does not destroy the argumentative function of the assertion. The argumentative move made is a move to the suggestion that the case in question is sufficiently like those cases the indefinite assertion summarises, i.e. the case in question is not a counter-example. This kind of argumentative move, however, is not one whose legitimacy can be properly understood to be deductive validity.

An analogous difficulty with pushing reflective legitimacy into a deductive mould concerns judicial reasoning from precedent. In Chapter III we saw that the jurisprudential Formalist insists on construing the ratio decidendi as a general, legal assertion. He is motivated by Deductivism: he needs universally quantified assertions as major premisses. But as we saw, rationes do not, and cannot serve that role. The reasoning which the judge engages in is not deductive, but reflective.

The conclusion to draw from these considerations is, it seems to me, clear enough: in some contexts and for special purposes it would be helpful and interesting to think of the legitimacy of nondemonstrative

reasoning as an approximation of deductive validity. One might, for example, wish to analyse inductive correctness as validity 'to some degree' (cf. Sellars). But if part of the story, it is only a small part. The problems and confusions created by this approach make it a non-starter.

Another approach to nondemonstrative legitimacy has also surfaced from time to time in earlier discussions. This approach centres around the claim that it is intuition, or direct apprehension, which informs us when arguments are legitimate and when not. Precisely what is intuited might be different things in different cases. For analogical arguments, as we have already noted, the closeness of the analogy, its success or failure in the context of the argument, or even the existence of an analogy, all of these might be what the intuitionist would claim we intuit. Or again, the intuitionist may want to say no more than that legitimacy or illegitimacy of arguments is intuited--we just see when some argument is legitimate, it is obvious, no more proof need be, or could be, forthcoming.

I have characterised the intuitionist approach as pessimistic because it is founded on the view that at some stage no further reasoning is possible, that at some stage we must resort to intuition to ground our faith in the legitimacy of nondemonstrative reasoning. But this is to see the resort to intuition as an unwelcome but unavoidable way of fending off the sceptic. It is true that those who opt for intuition often are very enthusiastic about this mysterious faculty. J. C. Hutcheson, writing on the nature of judicial reasoning, argues that such reasoning is neither deductive nor inductive, but he then goes on to extol judicial intuition:

Judges who have tried many patent cases, who have heard the testimony of experts, the one affirming the matter to be merely an advance in mechanical steps, the other to be invention of the highest order...know that for a just decision of such cases no objective criteria can be relied on. They well know that there must be in the trier something of the same imaginative response to an idea, something of that same flash of genius that there is in the inventor, which all great patent judges have had, that intuitive brilliance of the imagination, that luminous quality of the mind, that can give back, where there is invention, an answering flash for flash. (416.)

But there is something odd about this enthusiasm; for it is unclear what has been added to the description of what the judge must do in order to ascertain novelty or invention when 'intuitive brilliance of the imagination' has been added. That is, if the judge can give reasons in favour of the claim that in this case there is novelty, then why should intuition ever enter the picture? 'Intuitively, this is a novelty, an invention' is not itself a reason; reasons for thinking that this (intuitively if you like) is a novelty are reasons. What job does 'by intuition' do?

The most famous argument for the necessity for intuition in all reasoning is due to Sextus Empiricus (who employs the argument in Against the Logicians as evidence in support of his scepticism about reasoning). A. C. Ewing gives substantially the same argument, in modern dress, in his Ethics:

Suppose I argue A, ∴ B, ∴ C. Now the argument is invalid unless B does really follow from A, but how can I know that it does so? I may be able to interpolate an intermediate proposition D which itself follows from A and from which B follows, but this only puts the problem further back. I must know that D follows from A, and though I might perhaps be able to interpolate a further intermediate stage, I obviously cannot go on in this way ad infinitum. Sooner or later, and probably very soon indeed, I must come to some link between A and the next term in the inference which I can see immediately to hold without being able to prove this by further argument. We may take it then that, if we are to have any knowledge by inference, intuitive knowledge must occur...(137.)

There are two unstated premisses of this argument which I want to question. And by questioning, and finally rejecting these two premisses I will be discrediting the intuitionist approach to nondemonstrative legitimacy.

The first premiss was briefly mentioned in an earlier discussion of ethical intuitionism in Chapter I, and is this. The need for intuition is only felt if one is presupposing the deductive pattern of justification and proof. In Ewing's argument this presupposition surfaces at two places, first when he claims that we cannot go on in the way he suggests ad infinitum, and second when he claims that we need some way of seeing that the logical link holds without being able to prove that it does by further argument. The Deductivist insists both that justification must come to an end and that that end must consist of foundational claims which themselves do not require, and could not be given, further justification. The intuitionistic approach to nondemonstrative legitimacy presupposes this picture of justification, and Ewing's argument depends on it.

The second unstated premiss in Ewing's argument in effect sets the stage for the introduction of intuition. The premiss is the claim that we need a source of certain and self-evident, foundational claims which will guarantee the validity of valid inferences. Some inferences and some inferential links must be self-evidently valid for any to be valid. The source of this set of certain and self-evident foundation claims is intuition. In the case of nondemonstrative legitimacy, the story would presumably be similar. In order for any analogical argument to be legitimate, there must be some whose legitimacy is self-evident. Or, if the point at issue is whether something is of kind K, then if argument by means of the presentation of cases is to be possible, there must be

some instances of K which are self-evidently of kind K. To take an example, if the point at issue is whether what Jones did last Tuesday was courageous, then, the intuitionist will claim, if reasoning by example is going to be a legitimate procedure, there must be cases of courageous behaviour which are self-evidently cases of courageous behaviour.

These, then, are the two essential missing premisses in Ewing's argument, premisses which make the argument work. A case against both of them has already been set out in previous chapters. The case builds on the claim that there is a model of justification and proof which is not only non-deductive, but also is logically prior to the deductive model, in the sense that the deductive model presupposes it. In particular, my case against these two premisses is built on the denial of the existence of, or the need for, self-evident starting places, coupled with the claim that not every proof is finite. I will expand on these last two points in reverse order.

Ewing's argument, and, generally speaking, any argument for the necessity of unjustified or unjustifiable foundational claims, rests on the proposition that a genuine proof must consist of a finite number of steps the last of which is what is to be proved. In formal deductive logics, a proof is explicitly defined to be finite. But this conception of a proof is unnecessarily restrictive, although in the context of formal logics it has obvious meta-theoretical advantages (e.g., theorems about the completeness or consistency of formal calculi could not be demonstrated if proofs of infinite length were allowed). What is unnecessarily restrictive is a picture of proof which portrays the reasons adduced as 'links of a chain' rather than as reasons which

severally co-operate in favour of the conclusion (Wisdom (2), 203). Or again, one might think of a proof as a network of connections, "like a piece of cloth, the total strength of which will always be vastly superior to that of any single thread which enters into its warp and woof" (Perelman (1), 122).¹ Keeping these images in mind, and setting aside the model of a logical or mathematical proof, it becomes apparent that we should not be quick to disallow proofs which can go on ad infinitum. We should, of course, be wary of arguments which involve infinite regress, circularity, or mere repetition. But none of these kinds of disreputable arguments need be employed in the resolution of the problem which Ewing sets out for us.

As we have already seen, if the question of the legitimacy of a deductive inference is raised, it need not be answered by an appeal to the covering inference rule. If it is not, then we are more or less at the point where Ewing begins his argument. But at that point it is possible to give an infinite number, or better (given human finitude), an indefinite number of inferences which are parallel to the one we are worried about. Now, my contention is that the very fact that we could go on indefinitely--citing case after case where the reasoning is legitimate and analogous to the case troubling us--counts in favour not only of saying that here we have proof, but also of saying that here we have conclusive proof. Conclusive, but not deductive. But, it might be objected, don't we need intuition to see the legitimacy of each of these parallel cases? This brings us to the first of the two points mentioned above.

The intuitionist claims that we need certain and self-evident starting places, that we need cases of things of kind K which are self-evidently of kind K. But not only is this claim false, it doesn't

even solve the problem which the intuitionist wants to solve. It is a notorious objection to the various versions of intuitionism--and most commonly levelled against ethical intuitionism--that the intuitionist has no adequate reply to someone who simply denies that some claim is intuitively obvious. He has no adequate reply because on his view intuitively-sanctioned claims are not themselves susceptible to further proof. So the intuitionist cannot argue against the sceptic, the best he can do is to accuse the sceptic of having an undeveloped or untrained intuition. That rejoinder, as one might suspect, does not carry much weight. But even if intuition could solve the problem the intuitionist needs to solve, it would still be false that we always require certain and self-evident starting places. For an argument like Ewing's pre-supposes a deductive conception of justification and proof. Supplying another conception, and another model, nullifies the need for certain and self-evident starting points, and hence it nullifies the need for intuition. In terms of the reflective model of justification, where indefinitely long proofs are allowed, the intuitionist's problem, the gap he wants to fill, disappears. Moreover, the reflective model avoids the scepticism which the deductive cum intuitive model is susceptible to. For on the reflective model every claim can be the object of controversy; and so, if the sceptic rejects a clear case, or doubts the relevance of citing such a case, then he can be answered: reasoning, in the form of new examples, new parallel cases, goes on--indefinitely if need be.

The intuitionist at this point might retreat a bit and rest his claim that intuition is necessary on something like the following. Look at the law of non-contradiction (or some instance of it), look at a claim like 'It seems to me that there is a patch of red over there', look at a

clear case of murder, or again look at the claim that such a case is relevant to the reasoning one needs to go through in order to decide whether a particular case of killing is a case of murder--these are the sorts of claims which we know intuitively to be true. We could go on to reason about these claims if anyone were to question them--granted. But the fact that some instance of the law of non-contradiction, or some clear case of murder stands firm for us, is what I mean when I say that there is intuitive knowledge, a direct source of certainty.

If the intuitionist were to resort to this sort of stand, then we should indeed be sceptical about the work which '...is intuitively true' is supposed to be doing. For now the intuitionist is not claiming that there are some topoi for which no argument could be given. By retreating from that position he has taken away the need which intuition was to fill. If intuition is no longer needed to secure a certain and self-evident topos for an argument, and if to say that some claims are intuitively true and some cases intuitively paradigmatic is no more than to say that some claims are clearly true, some cases clearly paradigmatic, then intuition becomes, as Wittgenstein remarks, "an unnecessary shuffle" ((3), sec. 213). For 'clearly true' and 'clearly paradigmatic', or better still, 'true' and 'paradigmatic', are terms which tag claims and cases which we would normally take for granted, claims and cases which we would normally expect never to have to defend or explain. And that is the end of the matter. Suppose, for example, we are explaining to someone why the dinner bill was so high and we begin by saying, 'Well, four coffees, twenty cents each, that's eighty cents...' at which point our friend interrupts with: 'Hold on, four coffees, twenty cents each, eighty cents, I don't get it'. We have reason to pause in wonder and

and disbelief. We had no idea that we had to go down that far in order to explain why the dinner bill was so high: $4 \times 20 = 80$ is so clear, so elementary. But what more is added if we say that $4 \times 20 = 80$ (or $4 \times 2 = 8$, or for that matter, $1 + 1 = 2$) is intuitively obvious? If we (correctly) abandon the view that by 'intuitively obvious' we must mean 'self-evident, and I can give no reason for thinking it so', then 'intuitively obvious', 'by intuition' and the rest simply don't add anything; they are just unnecessary shuffles.²

Leaving accounts of nondemonstrative legitimacy which are based on the 'approximate validity' and the intuitionistic approaches, I want now to look at three far more interesting accounts. These accounts centre around propositions and images which point us in what I take to be the correct direction.

The first of these approaches is based on the plausible suggestion that reflective reasoning can be given a pragmatic vindication--in the manner in which some philosophers have argued induction can be vindicated. That is, the legitimacy of reflective arguments might be thought to be tied to, roughly speaking, their results. The best way of laying out the pragmatic approach is to try to get some idea of how to fill out the notion of 'results' used in this context, to see what can be said about the purpose served by reasoning generally. For this essential ingredient we can turn to the father of Pragmatism, C. S. Peirce, who, in his "The Criterion of Validity in Reasoning" writes:

Facts are hard things which do not consist in my thinking so and so, but stand unmoved by whatever you or I or any man or generations of men may opine about them. It is those facts that I want to know, so that I may avoid disappointments and disasters. Since they are bound to press upon me at last, let me know them as soon as possible, and prepare for them. That is, in the last analysis, my whole motive in reasoning. (Vol. II, section 173.)

The purpose of reasoning, then, is that of providing us with facts about the world, such facts as would affect our lives in possibly unhappy or disastrous ways were we to remain ignorant of them. (The assumption can easily be made here that every fact about the world is worth knowing, in the sense of being useful to us in some way or another.) The legitimacy of a piece of reasoning is thus tied to its success, where success is defined as and measured in terms of genuine epistemic advance. Or, to speak more strictly, the pragmatist may consider the question of legitimacy particularistically or methodologically--the first if he considers the legitimacy of each argument or inference to be a direct function of its success, the second if he considers the legitimacy of particular arguments or inferences to follow from the overall legitimacy of modes, schemes, patterns or general policies of reasoning. Peirce may have taken the particularistic approach--it is very difficult to tell--but the majority of those who have adopted the pragmatic approach to truth, legitimacy or justification have adopted the methodological approach. And there is a good reason for this, or at least it is a good reason in the context of the pragmatic approach: It is possible that a clearly fallacious argument might, by chance, inform us of a fact. But the pragmatist's interest in legitimacy is motivated by on-going concerns about future disappointments or disasters; so he is interested, not in the odd success, but in policies of reasoning which, if followed, would regularly yield genuine epistemic advance.

To see how this approach to the issue of reflective legitimacy would proceed, I want to consider how it operates in a wider epistemological context. I want to consider the pragmatic approach to justification as it fits into a coherentist epistemology, that is, a theory

of knowledge founded upon a coherence theory of truth. Coherentist epistemology is particularly appropriate here since most of the epistemological remarks which have been made up until this point, are, or at any rate sound like, remarks which coherentists frequently make. In particular, the coherentist would be sympathetic with remarks made against the priority of the Euclidean, or foundational model of knowledge and justification, with remarks about proof (e.g., the metaphors of Wisdom and Perelman which I have quoted and affirmed have a distinct coherentist flavour), and at least partially sympathetic with remarks made about the potential vulnerability of all claims. (Such affinities warrant the claim that the view of the nature of reasoning being developed here would be at home in a coherentist epistemology; they do not warrant the claim that what I am developing here is a coherentist epistemology. I am interested in certain ways of talking, certain images, not in adhering to a full-blown theory of knowledge.)

Methodological pragmatism is adopted as a procedure of justification in N. Rescher's coherentist epistemology presented in his The Coherence Theory of Truth (especially Chapter XIII and further developed in Rescher (2)). The justification which methodological pragmatism is to provide is, however, of a very special sort. On the coherentist account, knowledge is conceived to be a system of interrelated theses which are linked together by an interlacing network of deductive and non-deductive connections (as opposed to being related hierarchially in the axiom-to-theorem, deductive fashion). In theory, one starts with an incoherent jumble of theses (perceptual and memory 'data') which are mutually inconsistent; one then tries to sift the true from the false, not by matching theses directly with the world, but by constructing an

adequate and consistent system of theses. The adequacy of the system becomes a function of its internal coherence, of the fact that everything stands in a relation of mutual coordination with everything else. In Rescher's theory, a complex mechanism for assigning plausibility ratings to new 'input'-theses is developed which makes it possible for the coherentist to make sense of an ever-increasing body of knowledge which maintains its coherence: our knowledge does not grow by the mere accretion of facts, rather as each new piece of data is added to the system the entire system subtly changes as new connections are formed and others are abandoned. But then, justification takes on a double role. One can either be interested in internal justification--exposing the connections which one point in the network has with others--or one can be interested in external justification--determining whether the methods of organising the system are adequate to the task.

External justification, Rescher claims, is provided by the pragmatic approach. One justifies the method (in this case, the method of coherence analysis) by reference to practical criteria: "success in prediction and efficacy of control" ((1), 324). That is, one is interested in the method's "success in enabling us to canalize our actions and expectations" ((2), 706). Rescher's concern here is obviously with a very substantial methodology, but there is no reason why techniques of reasoning, say analogical reasoning, could not be subjected to the same sort of external justification. What would have to be said is something like this. A technique of reasoning is legitimate if (or, conceivably, to the degree that) use of that technique, in the long run, is pragmatically successful. Whether we take pragmatic success to consist in Peirce's avoidance of disappointment and disaster, or Rescher's efficacy in control or canalised actions and

expectations is unimportant. What is important is that this approach to the question of reflective legitimacy is not suitable for the bulk of reflective arguments. The pragmatic approach does have the advantage of putting deductive and non-deductive reasoning on all fours on the issue of legitimacy. But it is tailor-made for reasoning which is expectation-oriented, not concept-oriented. To succeed in reasoning about concepts is to come to a clearer understanding of those concepts; but that understanding is not always manifested in correct prediction.

We must say 'not always' here since there are borderline cases where the results of inductive reasoning produce greater conceptual understanding (see the discussion of von Wright's phosphorus example in Chapter Four). But, on the whole, success in reflective reasoning is not measurable by a direct comparison of conclusion with facts about the world. The problem is not that reflective reasoning is not answerable to external justification on pragmatic grounds. Not at all: a clearer understanding of a concept like negligence will have an effect on how we talk about negligent acts, and this information will canalise our actions and expectations. The problem is rather that there is no independent way of determining whether reflective reasoning succeeds in the pragmatist's sense. That is, the conclusion of a reflective argument is not a prediction, direct and independent confirmation of which is in principle forthcoming. The conclusion is more like a recommendation about how one ought to understand a concept. Like induction, reflective reasoning is ampliative, in that the conclusion goes beyond, and may be far richer in content than, the sum total of the premisses or reasons adduced in support of it. But while the conclusion of an inductive argument is invariably some fact or thesis which is (at least in principle) directly

and independently accessible, typically the conclusion of a reflective argument is a claim which expresses a new understanding of a particular case or a general concept. And that new understanding is only accessible by means of reflective reasoning, and can be affirmed or questioned only by means of further reflective argument.

This, then, is one reason why the legitimacy of a reflective argument can only be approached by reasoning in kind. Whether a particular argument counts as a good one is at bottom a question about what it means to reason, and what it means to be rational. It is a question about what it means for a particular topos, in a particular context, to be acceptable, and what it means for claims, examples and parallel cases to count as reasons for or against other claims.

But if the methodological pragmatist approach is ultimately inadequate in dealing with reflective legitimacy, might not the picture which the coherentist gives of in-system, or internal justification apply? That depends on how we are to understand coherence. On the level of the metaphors used and the image put forward, there are close similarities between Wisdom and Perelman and coherence theorists. Thus Rescher writes,

For the coherentist, knowledge is not a Baconian brick wall with block supporting block upon a solid foundation, rather an item of knowledge is like a node of a spider's web which is linked to others by thin strands of connection, each alone weak but all together adequate for its support. ((1), 319.)

But what is the logical character of these 'thin strands of connection'? For Rescher they are deductive and inductive connections (at least in (1); in (2) there is the suggestion that some of these connections might be non-deductive and non-inductive (see, e.g., (2), 699)). Yet, Rescher is very insistent about the point that 'truths of reason'--by which he

means logical axioms and deductive consequences of these--are not themselves nodes of the spider's web ((1), 361-365). But suppose we drop these two restrictions. Suppose, that is, we (i) characterise the connections very generally as 'good reason in favour of'-connections; (ii) return 'truths of reason' to the system; and finally, (iii) introduce into the system all claims of the form 'x is a good reason in favour of y'. These changes (and they are, admittedly, extensive changes) would obviously introduce into the coherentist theoretical picture a thorough-going fallibilism: not only are non-deductive, non-inductive connections open to criticism and support--not only are they what I want henceforth to call arguable claims--but so too are deductive and inductive connections, the axioms and first principles, the inference rules and rational policies.

The result would be a 'system' only in a Peckwickian sense; it certainly would not count as the sort of system which the coherentist theorist wants to construct (since the sort of system he has in mind must be consistent). Perhaps it would be better not to call this a system at all: it is, after all, a mere pool of possible topoi, a collection of possible starting points for arguments. Perhaps it would be possible to identify (in a manner which could only be described as 'holographically') consistent sub-systems which would match up with the closed and optimally determinate systems discussed above. But regardless of how the details would go, one aspect of the resulting picture is clear enough: The legitimacy of particular 'x is a good reason for y'-connection would be an arguable question, a question calling for reasoning. This point will be of considerable importance to later discussions.

The coherence view supplies a part of the story, but not all of it. Suppose we move on to another account which is as insightful as it is iconoclastic.

Another quote from Peirce can start us off. In one of his letters to Lady Welby, Peirce wrote: "I regard Logic as the Ethics of the Intellect--that is, in the sense in which Ethics is the science of the method of bringing Self-Control to bear to gain our Satisfactions" (Wiener, 415; and see also Vol. II, sections 7 and 52). The suggestion here is that logic, the study of reasoning, is at bottom a normative study, or in Peirce's terminology, a normative science. Peirce's reason for thinking this is simple enough, and it is very forceful: the science of reasoning teaches us how we ought to reason. Now we have already seen that Perelman (after Aristotle) identifies the fundamental rule of justice (in law and in morality) with the fundamental rule of reasoning--'treat like cases alike'. This is in the spirit of Peirce's remark. Indeed, the normative nature of rules of reasoning has impressed several philosophers. R. Edgley in his discussion of practical reason states flatly that "...matters of logic are themselves, explicitly or by implication, evaluative or normative..." (134). And J. O. Urmson in "Some Questions about Validity" has argued that deductive validity is not immune from the sceptic's demands for justification since to call an argument valid is not just to classify it, it is at least in part to evaluate it.³ But, as it stands, the claim that logic is entirely or partially normative does not immediately suggest an approach to the question of legitimacy. It is possible that Wisdom had in mind the normativeness of talk of reasoning when he took a quick poke at the moral sceptic with the suggestion: "suppose the goodness of a person were as objective as the goodness of an argument" ((1), 217). But how might we set about capturing the legitimacy of arguments in normative terms?

Perhaps the best place to begin is with the question, What is the

point of saying that the philosophical study of reasoning is normative? The point of this way of talking, it seems to me, brings up a methodological consideration which has important consequences in this context: a philosophical study of reasoning is not merely a descriptive study of how people reason (or what they take to be reasoning); rather it is that plus an account of how people ought to reason, an attempt to make sense of the concepts of reasoning and rationality. And as such, it is a normative account. So, adopting this way of talking, we can say that an argument is legitimate just in case it is what an argument ought to be like--it is a good argument. To say that someone is arguing legitimately is to say that he is arguing as he ought--he is arguing well. The objection that by speaking of 'good arguments' and 'arguing well' we are making it very difficult to distinguish respectable from disreputable forms of persuasion can be met by taking the normative mode seriously. We can say, following Aristotle, that a good argument is one which ought to persuade people, or, following Rawls (e.g. in (1)), one which would persuade a competent judge. There is, of course, a circularity here (one which is accentuated when the issue at hand is moral reasoning). Can this be avoided, or need it be?

In a recent account of the nature of moral reasoning, T. D. Perry has offered certain procedural requirements for being rational:

...we ought to take full account of the facts that are relevant to the question; and we should form our judgement or make our decision when we are in a psychologically normal state; and that judgement or decision should be disinterested, i.e. impartial and universalizable. (37.)

And one again, Rawls in his A Theory of Justice sets forth a very complex series of procedural rules and requirements which specifies a set of deliberators with certain (highly fictitious) properties, a setting and

indeed an entire argumentative environment which, he feels, would guarantee that autonomous reasoners would arrive at a just solution as to how society ought to be organised. (118-192). The point of these approaches is to meet, at least partially, the objection that the general normative approach to an understanding of rational arguing is circular. Perry goes further in trying to defuse the charge of circularity by adopting two views which are, in the context of moral reasoning, strange bedfellows: (i) the view that the normative procedural rules he suggests--taking full account of the relevant facts, being in a normal state, and being disinterested--are intuitively obvious (60); and (ii) the view that with moral reasoning we are not even trying to certify the truth or falsity of claims--since moral assertions are only relatively true or false--rather we are trying to justify them (which, Perry claims, is something different). The tension in Perry's account is felt particularly strongly on the question of whether substantive moral claims follow from his procedural rules--whether, that is, any substantive claims are excluded by the procedural rules. (In fact he explicitly contradicts himself at various points on this question--vide 60 and fn. 7, 189-190). But if we take the normative character of rules seriously, if we agree that rules which determine what is to count as being reasonable are themselves substantive moral assertions, then this tension can be avoided. But can we avoid circularity?

It seems to me that an objectionable circularity in the normative approach to legitimacy cannot be avoided, if, that is, we are wedded to Deductivism. For if our picture of what it is to reason is limited to a picture of deductively conclusive and compulsive inferential links and self-evident (or otherwise unarguable) topoi, then the claim that a legitimate argument is one which ought to persuade (or would persuade

a reasonable person) is caught in the tightest of circles, a definitional circle. For if the normative approach is taken seriously along Deductivist lines, there would be a need for a set of strict necessary and sufficient, normative conditions for what it is for an argument to be legitimate. At least some of these conditions would, given the nature of the approach, make reference to a person--or some set of people, real or hypothetical--who is (or are) reasonable. A legitimate argument would thus be one which ought to persuade, that is, one which a reasonable person would be persuaded by. But, a reasonable person would have to be, inter alia, the sort of person who was able to identify legitimate arguments. (Analogous difficulties have arisen in the case of the 'Ideal Observer' theories of ethics, theories which define the right thing to do in terms of an Ideal Observer who is in turn defined as possessing certain ideal properties, such properties as would enable him to properly identify the right thing to do when confronted by it.)

But, if we drop the Deductivist prejudice in favour of compulsive inferential links and self-evident or otherwise unarguable topoi, are we still left with an objectional circularity? We are left with a circularity, but not an objectionable one. To show why the circularity is not objectionable, I want to consider one last approach to the question of legitimacy which incorporates the normative approach.

The approach I have in mind is a conceptual approach, in a word it is the approach which builds on the claim that the concept of argumentative legitimacy follows from what it means to reason, to be rational. The 'follows from' here is the crucial parameter of the approach. The connection these words refer to is obviously an a priori, conceptual connection; but more needs to be said about it.

I want to begin with a version of the conceptual approach which should be discounted. This version reads 'follows from' as 'deductively follows from'. On this reading, the conceptual approach can be paraphrased in this way: criteria for the legitimacy of forms of reasoning are deductive consequences of criteria for what it means to reason and to be rational. It should be noted that the concept of legitimacy which this approach seeks to account for is just that concept which figures in our second requirement of reasoning.

For this approach to be successful, it must be the case that the concepts of reasoning and being rational are, or could without generating difficulties be made to be, determinate. As it stands, this approach is open to the obvious objection that it only serves to shift the problem from understanding the nature of legitimacy to understanding what it means to reason and to be rational. But, there is no obvious reason to think that these latter concepts are any clearer than the former one. To meet this objection, it must be shown that we have a firm grip on what it means to reason and to be rational, the sort of grip necessary and sufficient conditions provide. Consider an analogous problem, the problem of understanding what it means for some set of evidence statements, reports of the results of experiments, and so on, to confirm a hypothesis. It has been suggested that we can understand the concept of confirmation once we understand what it means for a hypothesis to be scientifically acceptable, since it seems plausible that of two rival hypotheses, one is more acceptable than the other only if the available evidence confirms one more than the other. But then, in order for this shift from confirmation to acceptability to come to our aid, the concept of scientific acceptability must be less problematic than that of

confirmation. But it could easily be objected that as a matter of fact acceptability is, if not more problematic, then at least as problematic as confirmation. Much the same story can be told here. For this version of the conceptual approach to come to our aid, the concepts of reasoning and being rational must be less problematic than that of legitimacy. But more importantly, if, as this version of the conceptual approach suggests, there are strict, deductive connections between these concepts, then we require the concepts of reasoning and being rational to be determinate.

From what has already been said in previous chapters, it is clear that the only way in which these concepts would be determinate is if what it means to reason were restricted to what it means to reason deductively. To be sure, if we do make this restriction, then at least some of the problems involved with legitimacy would be solved. If what it means to reason is just to follow the rules of deductive logic, and if what it means to be rational is just to accept deductive consequences of propositions already accepted, then legitimacy is just deductive validity. But enough has been said to show why such a view of legitimacy is unacceptable. It is not that deductive reasoning is not reasoning, or that being deductively rational is not being rational. Rather, deductive reasoning is not all there is to reasoning, and accepting deductive consequences is not all there is to being rational.

I want, then, to claim that the concepts of reasoning and being rational are nondeterminate. But if this is so, then the foundation for this version of the conceptual approach collapses. As I have tried to show above, deductive inferential links require, impose or presuppose the proper sort of conceptual environment, namely a determinate one.

Without that environment, the scope of possible deductive relations which are informative is severely limited. But it is the keystone of this version of the conceptual approach that there are meaningful deductive connections between the concepts of reasoning and being reasonable, on the one hand, and legitimacy on the other. This version will not do the job.

At this point it might be objected that I have burned too many bridges behind me. I seem to be facing a dilemma: either what I will be forced to say about the nature of legitimacy will rely on such weak requirements that any sequence of claims could be argued to constitute a legitimate argument and count as reasoning, or I will be forced to admit that what we need now is a strong dose of the Deductivist's no-nonsense approach to legitimacy, firming up explicit criteria and constructing sets of necessary and sufficient conditions for legitimacy. This dilemma incorporates the three pessimistic attitudes which have been following us around. Either we should be sceptical about the legitimacy of reasoning, or say that the whole matter should be left to our intuitions, or we should take the bull by the horns and try to fit reasoning as a whole into the deductive mould. Once again, I want to avoid all three pessimistic attitudes. And I want to get between the horns of the dilemma I seem to be facing. To do these things I want to return to the conceptual approach and show both how it can do the job for us and why it vindicates Wisdom's claim that in order to understand when and how arguments are legitimate we must consider the cases.

I return to the conceptual approach because there is little doubt that it is correct. In order to understand what it means for a piece of reasoning to be legitimate we must bring in considerations both of

what it means for something to be a piece of reasoning and what it means to be rational. But, do we have a clear picture of these concepts, or is this approach just doomed as an explanation obscurum per obscurius? It seems to me that we have a clear enough picture of these concepts, if we draw together the various elements of the picture which have already been laid out.

The principal thesis which was developed in earlier chapters and brought to this one was that, in the end, all reasoning is reflective case-by-case reasoning. Deduction and induction were argued to be specialised forms of reasoning (presupposing an appropriate conceptual environment) which depend for their legitimacy on the legitimacy of reflective reasoning. Reflective reasoning, it was argued, is reasoning over questions about concepts, it is conceptual reasoning. It is also ampliative reasoning, reasoning which is--to use Wisdom's helpful term--'horizontal' rather than vertical, reasoning which forms an interconnected web of reasons which severally cooperate in support of a conclusion. Moreover, as we have discovered in this chapter, reflective reasoning is argumentative, proceeding as it does from arguable starting points and continuing, if need be, indefinitely. It is reasoning which proceeds towards greater conceptual understanding by means of a consideration of examples, parallels, analogies, clear and borderline cases. It is reasoning the appropriate model for which is the legal forum, not the calculating machine. It is, finally, reasoning which brings to our attention reasons for looking at the case or kind of case at issue in one way rather than another.

Now there are conceptual connections between this picture of reasoning and the concept of legitimacy which are vital and need to be

explored. One of these has to do with an issue which was set out in my earlier treatment of rhetoric, and can be put in this way: Do we lose the distinction between respectable and disreputable reasoning if we conceptually tie the legitimacy of arguments with an argument's capacity to persuade?

The worry expressed in this question is that legitimacy should not be linked with the mere capacity to evince agreement. This Thrasymachian-like stand is surely objectionable. But it can be avoided by relying on our non-deductive model of reasoning. And when we do rely on this model we invariably end up making normative claims: Legitimate arguments are those which are worthy of persuading people, arguments which would persuade rational people, people who are capable of reasoning and are concerned to discover what is true and just--competent judges, phronimoi, and perhaps too Lord Denning's "right-minded" people (4-8). By talking in this way we are in effect insisting on a distinction between techniques of reasoning which allow rational people to apprehend what is true or most likely from techniques which merely induce adherence. But we are also bringing into the picture what it means to be rational; we are bringing out the normative quality of that concept. So, suppose we shift our attention to the concept of being rational.

What I want to suggest is that in order to be rational one must, in so far as one is capable of doing so, satisfy one's argumentative responsibilities. That one has these responsibilities as an arguer is a result, I claim, of what it means to be rational. Some of these responsibilities have already been mentioned, they are T. D. Perry's procedural requirements: one must take full account of the relevant facts, be disinterested and in a psychologically normal state. These

only scratch the surface, however. There are other, more important, argumentative responsibilities. One need only look through the Platonic Dialogues to find them. In the Dialogues, Socrates time after time scolds his interlocutors when they fail in their rational role, when they fail to acknowledge the importance of a Socratic counter-example, refuse to give reasons for their views, refuse to cite examples of general claims they want to hold, or disrupt the continuity of the argument. One simply has no right to refuse to give reasons for a view one is claiming to be true or to refuse to respond to reasons given against it, if one is also claiming to be rational. It is, of course, a different matter if someone cannot see the point of a particular argumentative move, cannot see that something is a counter-example to a claim he has made, for example. But such a person is not shirking his responsibilities as a rational arguer. But refusing to acknowledge the possibility of counter-example is removing oneself from rational argument, breaking a rational rule, shirking an argumentative responsibility.

How this feature of rationality affects the question of legitimacy is clear. I have been maintaining that the question whether in a particular context some argument is legitimate is itself an arguable question, a question which can be meaningfully raised and for which other sorts of reasons and arguments are appropriate. (It should be noted that by 'arguable' I do not just mean contentious or moot. To be sure, all contentious or debatable issues are arguable issues. But not all arguable issues are necessarily contentious. For example, it is my claim that even the principle of non-contradiction is arguable, in the sense that--in the unlikely event that someone would take exception to it--reasons can be brought to bear on the truth of the principle. And

these reasons will take the form of particular instances of the principle, clear cases which show its truth.) But then, if the legitimacy of a particular argument is itself arguable, then one has certain argumentative responsibilities with regard to the arguable issue of legitimacy. In particular, (i) one does not have the right simply to refuse to acknowledge the legitimacy of an argument--one must be prepared to give reasons for thinking that the argument is illegitimate; and (ii) one does not have the right to refuse to consider the reasons given for thinking that one of one's arguments is illegitimate--one must be prepared to meet the challenge, either by arguing that the reasons given are inadequate, or by showing that there are other, and better, reasons for thinking that the argument is legitimate. And these responsibilities are part of what it means to be rational, to proceed in a reasonable manner.

Now, since the question of legitimacy is a question both about the topos and the scheme or pattern of reasoning employed, the reasons that are brought up for or against the legitimacy of an argument must be relevant to one or the other (or both) of these aspects of an argument. And the sorts of reasons that can be given one way or another will vary with the kind of argument one is considering. If the argument is intended to be purely deductive, then a good reason for thinking that it is illegitimate is that it is formally invalid, a good reason for thinking it is legitimate is that it is valid. These are the obvious reasons to give. Equally obvious are reasons that pertain to the topos of the argument, reasons that have to do with the truth, plausibility or relevance of the premisses which are used.

The Deductivist is forced to claim that nothing more can be said for or against the legitimacy of an argument beyond what can be said about

its formal validity or its soundness. Against this I want to claim that a great deal more can be said, that reasoning about legitimacy can go on. There are things that can be said if someone fails to appreciate the force of the claim that an argument is formally valid. But at this point we must rely on case-by-case reflection and give examples of deductively valid or invalid arguments which are structurally identical to the argument in question. These examples count as reasons for a deductive argument's legitimacy or illegitimacy; they also count as reasons why it is that formal validity and invalidity are relevant to the question of legitimacy. On the other hand, if the question is that of soundness, then reasoning can also go on in the face of challenge. Here we must cite cases or features of cases which are parallel to those cited in the premisses and which are relevant or irrelevant to the issue which the argument is being addressed to. In the end, the arguable issue of legitimacy, like other questions which require reasoning, can be resolved by appealing to non-deductive, non-inductive reasoning.

The point here is that the issue of legitimacy should not be viewed in the Deductivist manner as an issue susceptible to a mechanical decision procedure. Rather, it should be viewed as an arguable issue, an issue which can be argued and for which reasons one way or another can be given. Reliance on a mechanical procedure in the case of deductive arguments is not irrational; but it is rational only because reasons for thinking that the procedure will not lead us astray can be given. Sound and valid deductions are perfectly legitimate; but it is rational to think that they are only because the question of their legitimacy is arguable. Appealing to principles is a legitimate way of reasoning; but it is rational to think so only because the legitimacy

of that way of reasoning can be rationally defended.

But, if we view the issue of legitimacy as an arguable issue, we soon see why considerations of the form of arguments--the rule, pattern or scheme of argumentation which an argument exemplifies--are relevant to the issue of legitimacy. These considerations are relevant because they serve to bring to the discussion an indefinite range of similar arguments which could be compared or contrasted with the argument we are considering. To say that an argument is an instance of modus ponens, of induction by enumeration, or an instance of one or more of the several dozen argumentative schemes which Perelman and Olbrechts-Tyteca discuss, is to draw our attention to infinite families of arguments of which the argument in question is a member. Of course, to say that an argument is an instance of modus ponens, and so is valid, is nearly always the end of the matter (unless there are questions to be raised concerning soundness). We are very impatient with those who would claim to fully appreciate what it means to say that an argument is an instance of modus ponens but who also claim that that doesn't settle the question of legitimacy. But, just like the fellow who had trouble with $4 \times 20 = 80$, the fellow who objects to the legitimacy of a modus ponens argument is not making a claim which is beyond the scope of reasoning. There are in principle an infinite number of reasons in favour of the claim that such an argument is legitimate; namely, the infinite number of examples of modus ponens arguments which could be brought to his attention. On the other hand, to say that an argument is analogical is nearly always not the end of the matter. But here too, the path we must travel in order to argue for the legitimacy of the argument is linked with examples.

So, considerations of the form of a piece of reasoning are always

relevant to the question of legitimacy, but these considerations carry different weight in different cases. The reason for this brings us back to a point developed in some detail earlier: one reasons differently in different conceptual environments, against the background of different sorts of systems--determinate, partially determinate or nondeterminate. The question of legitimacy, then, is relative to both the kind of argument being considered and the context in which it is placed. As we have said, legitimacy is a matter both of the topos and the scheme; but it is the former which sets the context, the argumentative environment, since it is the former which either implicitly or explicitly lays out what, for the purposes of the argument, is not in question. In the court of law the context is fairly rigorously set out by means of presumptive and procedural rules--rules of the burden and onus of proof, for example. In other contexts, sometimes more and sometimes less is wrapped up in the topos. A clue as to what is part of the topos of an argument comes from a consideration of what is relevant to the argument. It would be pedantic--because contextually irrelevant--to bring up questions about the existence of other minds when the issue is whether Jones is justified in being furious about Smith's refusal to support Jones's bid for MP. It would be out of place to bring up worries about the ontological status of numbers in a demonstration that every Cauchy sequence of rational numbers converges. And so on. At the same time, however, there are no iron-clad rules for what is or what is not relevant to any arguable issue. Arguments grow and shift ground, they move off on tangents some of which prove on closer inspection to be more relevant to the issue than they did at first. Sometimes arguments concentrate on very specific questions, glossing over related major issues; sometimes the major issues are debated

while the details are assumed, for the purposes of the argument, to be irrelevant. In short, the question of relevance can only be approached on a case-by-case basis.

But given the context-relative nature of arguments, Wisdom's three "philosophical operators" (vide (3))--'You might as well say', 'Exactly so' and 'But this is different'--perhaps capture the moves and patterns of argumentation better than either formal deductive or inductive patterns, or the informal schemes found in The New Rhetoric. For reasoning is fundamentally argumentative in nature, involving an exchange of reasons for and against, a dialogue. Formal deductive reasoning is notoriously one-dimensional, it is a monologue. Socrates's objections to long speeches and his preference for a question-and-answer style reflects an important insight which undue emphasis on mathematical reasoning obscures: one must always keep in mind the topos of an argument, one must always try to expose the hidden premisses, the background, the presuppositions. Wisdom's philosophical operators are argumentative operators; they serve to expose both the consequences and the presuppositions of claims and arguments. If we become, as modern logicians have become, overawed by the formal structure of arguments, if we become enchanted by the pattern and quickly pass over the content and the topos of an argument, then we are likely to be lead astray by bad arguments. And this is especially true for arguments in those areas where the presuppositions cannot be easily laid out from the start and agreed upon, that is, those areas where we seem to have the most problem getting off the ground. With philosophical queries and problems, as Wisdom has remarked, more often "it's not the stuff, it's the style that stupefies" ((1), 28). It's not that we lack data, it's rather that we

are mislead by the style with which philosophers mould that data into theoretical claims about what it is to know something, or what it is for a table to be solid, a room empty or a man evil.

But Wisdom's philosophical operators are also perfectly at home with the question of the legitimacy of an argument or a way of reasoning. Here we are interested in the consequences of calling this argument, or this way of reasoning, legitimate; here we are interested in the concept of the legitimacy of reasoning. Suppose someone were to argue for the moral permissibility of abortion by making an analogy between a human fetus and a parasitical organism. They are both wholly dependant on another for their continued existence, they both to one degree or another affect the health and well-being of another, and so on. We could object: You might as well say that a one-month old child, or a ninety-year old invalid, or for that matter a temporarily bed-ridden adult are all parasites. This move brings out a reason for thinking that the original analogy is a bad one, one which misrepresents the status of a human fetus by exaggerating one feature (from dependence to parasitical existence) while obscuring another feature (the potential personhood of the fetus). Or suppose someone were to argue that if we were to take seriously the view that animals have a prima facie right to life, then in the case where a human baby is attacked by a rat we should not be able to decide which of the two ought to live. We could object: But this case is different. Granting animals a prima facie right to life means that we should not kill them gratuitously-- in this case there is a good reason for trying to save the baby's life at the expense of the rat's. Or again, should laboratory animals be subjected to painful experiments, should they be killed for the sake of

scientific advancement, if they have a prima facie right to life? Perhaps this case is different too; or better, perhaps it depends on the particular circumstances: Is the experiment an important one, would it truly advance our understanding of a disease like cancer, or is the experiment merely a trivial exercise, a wanton waste of animal life and an unnecessary production of animal suffering? There are cases, and there are cases.

But lastly, sometimes the argument hits the nail on the head, sometimes the analogy employed or the borderline case offered works successfully and brings into prominence a feature of a disputed case which needed to be brought out. One might think here of Judge Macnaghten's instructions to the jury in the important abortion case of Rex v. Bourne ((1939) 1 K.B. 687). Macnaghten brought to the jury's attention the fact that although previous cases of abortion were judged to be legal on the grounds that the phrase 'danger to life' (in the Offences Against the Person Act) clearly applied to an instance of direct physical threat to the life of the mother, still the phrase may also apply to cases of severe and prolonged deterioration of mental health. The analogy drawn, the similarities exposed, between physical threat and mental threat should, in this context and in light of the circumstances of the case before Judge Macnaghten, call up the application of the last Wisdomian operator: 'Exactly so'.

In each of these cases we are making a conceptual comment about the legitimacy of an argument. In the first two cases we are questioning the legitimacy of the argument, not by applying the relatively ineffectual tool of deductive validity/invalidity, but by arguing either against the starting point of the argument (the value of the analogy) or the consequences which might be drawn from that starting point (what one would be committed to if one viewed the analogy as being based on relevant

similarities). In the last case, we are noting that the argument is founded on a conceptual question--what constitutes a 'danger to life'--and an interpretative point which ought to be brought out since it illuminates the rationale and the spirit of a part of a statute.

This suggests that legitimacy is both an arguable issue and one which does not always admit of a conclusive, Yes or No, answer. And this I think is true; legitimacy is, for the most part, a matter of degree. To be sure, legitimacy is less a matter of degree in the case of deductive arguments. But it is still a matter of degree even there: validity is a Yes or No matter, but the epistemological status of the premisses need not be. Indeed, when deductive arguments occur outside of the context of a formal calculus, soundness is nearly always a matter of degree. Most of the reasoning we engage in outside of mathematics, formal logic and highly systematised areas of science concerns the status of potential premisses for relatively unimportant and usually very uninteresting deductive inferences. (It is significant that the deductive complexity of reflectively complex arguments in the law, morality, philosophy and elsewhere rarely exceeds that of modus ponens. Usually when deductive moves become possible all of the difficult and important reasoning has already been completed.)

It is also important to note that a substantial part of the reasoning we engage in outside of determinate systems is reasoning about the legitimacy of other pieces of reasoning. Arguments are not like individual bricks which can be stacked one upon the other to form a solid wall; they are rather like individual threads in a fabric; each may be challenged and either survive or be rejected, but if there is a case to be made, the whole, the fabric will remain intact. Arguments over

important and complex issues are constantly being taken apart and rebuilt, sometimes strengthened, sometimes weakened.

We return, then, full circle to Wisdom's claim that the issue of legitimacy is itself a matter for reflective reasoning: to see whether a particular argument, or a way of reasoning is legitimate one must look at the argument or way of reasoning; one must consider the cases of good and bad reasoning, seeing how we can go wrong and seeing when we have got it right. I have tried to supply some of the theoretical background to Wisdom's claim, to put his claim in the proper setting--in short, to make a case in favour of his claim. This setting is constructed out of the basic recognition that the issue of legitimacy of reasoning is a conceptual issue which is linked to the concepts of reasoning and being rational. These latter concepts, moreover, contain a normative element, namely the fact that there are certain responsibilities which reasoners incur. This in turn underscores an aspect of the question of legitimacy which was implicit in earlier discussions: the issue of legitimacy itself is arguable. But then, reasoners have responsibilities with regard to the question of the legitimacy of the arguments they bring forward: they must respond to the challenge that their arguments are illegitimate, and they must respond by giving reasons for thinking that they are legitimate; in short, they must argue for the legitimacy of their arguments.

Defending deductive arguments by marching axioms and inference rules past one's objector's nose is a defence. But it is a defence which can be challenged. And when it is, deductive arguments, like non-deductive arguments require non-deductive, reflective justification. But given that a reasoner must defend himself against objections, must

rationally withstand criticism, not every sequence of claims will count as legitimate reasoning. For a particular argument it may not be the case that there is a single, conclusive reason for thinking it illegitimate. But if the argument is illegitimate, that fact will show itself as the individually inconclusive, but severally conclusive reasons adducible against it build into an interlocking collection of arguments, a conclusive case, which only an irrational man could reject.

But what of the objection that Wisdom's answer is circular? Since the question of the legitimacy of reflective reasoning--which is to say, the legitimacy of all reasoning--is a conceptual question, linked in the ways I have suggested to other concepts, we have a way of meeting this objection. Any investigation into any concept is circular in just the sense in which Wisdom's solution is circular. One must begin somewhere; and one begins with what one is confident about. Distinguishing good and bad arguments is much like distinguishing sane from insane men. One could conceivably consider every human being alive, perhaps every human being who has ever lived, and ask of each, 'Sane...or insane?'. The question presupposes that we have some idea, however inexact and however tentative, of what could count as an answer to the question. We would undoubtedly change our lists of the sane and insane in light of new cases--borderline cases for example. Our paradigms would change and the borderline itself would shift. But, as we proceed our understanding of what counts as sane and what counts as insane will increase and our nosography will be revised accordingly. We will come to have a better idea of what to look for--the tell-tale signs, the traumas in childhood, the differences between genuine eccentricity and genuine mental ill health. We will come to see more clearly why we have the distinction

in the first place, what difference the concepts point our attention to. Our case-by-case reflection will produce conceptual clarity. Wisdom's answer to the question of how we can distinguish good from bad arguments is based on the same, conceptual point: in so far as the question makes sense to us, we have some idea of what to look for. Moreover, we have a good idea of how to proceed: we look and see.

Footnotes

Chapter V: Reflective Reasoning

(Page 269)

¹It is interesting to compare Wisdom's and Perelman's two images with a similar image offered by Wittgenstein. Speaking of our understanding of concepts, Wittgenstein notes that we extend our understanding "as in spinning a thread we twist fibre on fibre. And the strength of the thread does not reside in the fact that some one fibre runs through its whole length, but in the overlapping of many fibres" ((3), section 67). The similarity of image here discloses another important feature of proofs: to prove something is to come to understand one or more concepts, but to understand these concepts without necessarily relying on a set of necessary and sufficient conditions (fibres running the whole length, threads making up the total strength) for the application of the concept. And compare the following characterisation of judicial proof on circumstantial evidence given by Pollack, C.B. (R. v. Exall (1866), 4 F. & F., 922, at 929):

It has been said that circumstantial evidence is to be considered as a chain, and each piece of evidence as a link in the chain, but that is not so, for then, if any one link breaks, the chain would fall. It is more like the case of a rope comprised of several cords. One strand of the cord might be insufficient to sustain the weight, but three stranded together may be quite of sufficient strength. Thus it may be in circumstantial evidence--there may be a combination of circumstances no one of which would raise a reasonable conviction or more than a mere suspicion; but the three taken together may create a conclusion of guilt with as much certainty as human affairs can require or admit of.

(Page 272)

²This is not to say that we cannot use these expressions, or that we are making some sort of conceptual blunder when we do use them. The intuitionist has gone wrong in supposing that this family of expressions adds something, that talk of intuition explains or accounts for the legitimacy of modes of reasoning or the fact that certain propositions, once we know what they mean, stand firm for us. Knowing something intuitively is knowing something without having to deliberate or reason about it. But if few have to deliberate over an elementary mathematical equation, it does not follow that no one could; nor does it follow that elementary mathematical equation cannot be supported by reasons.

(Page 279)

³Urmson's treatment of deductive validity is, by the way, another example of how the general Deductivist approach to justification fails to counter scepticism. Urmson makes the typical Deductivist move from 'x is an evaluative term' to 'we require principles to determine whether,

for any particular example, it is x or not'. The result is, predictably, a regress. For if, as Urmson thinks, the question 'What good reasons can be given for rating paradigm examples of valid arguments positively?' can only be answered by providing standards or criteria for rating positively, then the immediate question becomes--once these criteria have been supplied--'Why are these the appropriate criteria?' To avoid this Deductivist regress, we must stay at the level of the examples: paradigms are paradigms because they are clear cases, and clear cases are such because we can proceed indefinitely in support of them.

CHAPTER VI

QUESTIONS REQUIRING REFLECTION

In the opening pages of Ethica Nicomachea, Aristotle advises that in a discussion of ethics we should be content with statements which are only for the most part true and arguments which are nondemonstrative, since

...it is the mark of an educated man to look for precision in each class of things just as far as the nature of the subject admits; it is evidently equally foolish to accept probable reasoning from a mathematician and to demand from a rhetorician demonstrative proof. (1094b 24-27.)

On the face of it, Aristotle's advice seems platitudinous: one surely shouldn't use a microscope to study social organisations, nor should one investigate the properties of tachyons by conducting surveys. There is that advice, but there is much more. It seems to me that Aristotle is saying here that there are different but related ways of being precise and different but related ways of reasoning which are appropriate for different kinds of questions and subject-matters. He is also saying, it seems to me, that there are different but related ways of being rational.

It is, however, easy to misinterpret this piece of advice, and to misinterpret it in one of two ways: (1) in ethics the best we can hope for are dubious claims and specious arguments (mere persuasion)--but we should be content with this, for this is what it means to be 'rational' in ethics; or (2) there is no way of comparing the claims and ways of

reasoning appropriate to one subject-matter with those appropriate to another--we must be content with the fact that there is no limit to the number of ways of reasoning, of proving, or of being rational.

Both of these interpretations miss the mark, and for the same reason. Both suggest that Aristotle is advising us to stretch the concept of rationality so that it will be correct to say that one can proceed rationally in ethics, despite the 'for the most part' claims and the nondemonstrative arguments. But Aristotle is neither saying nor insinuating that we should introduce a sham sense of 'rational', or that as far as being rational is concerned, anything goes. Rather he is suggesting that there are important systematic differences between mathematics, logic, physics on the one hand, and ethics, politics, rhetoric on the other. And because of these differences, there are different, but equally rational, ways of reasoning which are appropriate.

I want to subscribe to this view. But I also want to fill it out a bit by putting it in terms of what I have said above. In this concluding chapter, I want to draw out some of the consequences of the view I have been arguing for in this thesis, namely that there is a way of reasoning--not a new logic, or process of reasoning, but a way of arguing rationally--which is neither deductive nor inductive but which is an eminently rational way of proceeding towards the resolution of certain kinds of questions. Making a case for the rationality of reflective reasoning from particulars, from examples and parallels, has been one of the purposes of this thesis. Yet, another purpose has been to try to give the formal logician his due without thereby acquiescing to the dogma of Deductivism. So, I want to begin by saying that while I would totally reject the claim that we never reason in a purely deductive or

inductive fashion, or that there is something wrong with deduction or induction, at the same time I want to join with Perelman in saying that

A careful study of the reasoning employed by the creative and original thinkers, both in science and in philosophy, would reveal that that reasoning is infinitely more varied than anything to be found in the manuals of logic or scientific methodology. ((1), 94-95.)

(Actually, I would prefer to say 'indefinitely more varied'--the realm of possible ways of reasoning cannot be as extensive as Perelman thinks.)

So, let me try to give the formal logician his due. One invariably reads in textbooks of formal logic that the nature of formal logic is such that the legitimacy of arguments constructible in the language is removed from the realm of controversy. In propositional logic, for example, proof procedures are wholly algorithmic: there are well-known, unproblematic, mechanical ways of determining whether some expression in the language is provable and whether the proof has been correctly constructed. And, given that a formal proof is the final arbiter of the formal validity of an argument, formally expressed, there is an important sense in which it is true that the question of legitimacy in formal systems is not arguable. The sense in which it is true is this. Formal logic comes complete with a highly systematised topos; indeed we might even say that it is nothing but an optimally determinate starting place. The system is identifiable by its vocabulary, formation and transformation rules and its axioms--these structural elements determine, explicitly and exhaustively, the criteria for the validity of arguments expressible in the language. In so far as one understands the structural components of such a system, one understands what it is for a proof to be correct.

But, the sense in which it is false that the question of legitimacy in formal systems is not arguable is that the structural ingredients

themselves are challengeable. Of course, they are rarely challenged, since we rarely feel the need to question or to try to argue for, say, the axioms of a formal system. We rarely have reason to be dubious about the structure of simple formal systems. (On the other hand, there are on-going controversies about the structure of more complicated and philosophically more contentious formal systems like modal logic, temporal logic, epistemic logic and the rest.) But even though we rarely have reason to question the topos of propositional logic, the axioms and rules can be challenged; reasoning is possible at this level. It may be foolish, even paradoxical, to challenge or question the law of contradiction; but if challenged or questioned, reasoning can find a foothold: instances of the law can be brought forward in support of the law. We have seen that an inference rule like modus ponens is arguable, and we have seen how the argument would go.

Deductive reasoning is not restricted to formal systems of logic, mathematics or natural sciences; it is ubiquitous. I have made a point of insisting that the proper home of deductive reasoning is within a system of one sort or another which is, or is assumed for the purposes of the argument to be, determinate. Deductive reasoning presupposes conceptual decisions; it presupposes predicate univocality. This is the sort of conceptual environment one finds within formal systems; but one finds it in everyday, informal contexts as well. In everyday situations, when we argue deductively, we usually are not troubled about the nature of the concepts employed--we are all too willing to assume conceptual determinateness, it makes for simpler and more forceful arguments. And most of our everyday deductive reasoning employs concepts, rendered by mass- or count-nouns, which name the familiar objects of our world--tables,

chairs, dogs, mortises and marrows. We tend to assume for the sake of simplicity that the application of these concepts poses no real problem (or if we are aware that the concepts we are using have fuzzy boundaries, we tend to ignore the problematic, borderline cases for the purposes of informal argument and discussion). It would be plausible to argue that everyday deductive reasoning takes place in a system, formed by these concepts, which is as familiar to us as the tables and chairs themselves, a system of everyday concepts. It might be infelicitous to use the word 'system' here--perhaps it would be better to speak of the world of common sense--but one thing is clear: we do not always reason reflectively because we do not always need to.

But sometimes we do need to, sometimes questions arise which our common sense cannot answer one way or another. Sometimes the question at issue is a question about a concept and the question generates a controversy over the application of that concept to a particular, perhaps borderline, case. For such questions and controversies, as I have argued above, deduction and induction are usually ineffectual: these questions and controversies call for reflection, a consideration of examples and parallels, of clear cases and of not so clear cases. In the end, whether a question or puzzle calls for deduction, for induction, or for non-deductive, non-inductive reflection is a question which can only be answered by looking at the cases. But there are clear cases where reflection is called for.

Wisdom at one place calls them 'riddle questions'--Is a tomato a fruit or a vegetable? ((1), 142): Can one lie without deceiving? Can one murder someone who asks to be killed? Can one attempt to murder someone who is already dead? Can one steal something which one already

owns? Is a human fetus a person or merely a part of the mother's body? Can a non-human animal have a right? Can one be lonely and yet have many friends? Was Joan of Arc truly religious or, e.g., pathological? Is someone speaking French if he is merely repeating sounds he has memorised? If Jones jocularly states 'There's strychnine in the leek soup you know' is he giving a warning if there is? Can one play chess without the queen? (Wisdom (2), 88.)

In the first chapter I suggested that questions like these possess what seems to be a perplexing conjunction of properties: We have no reason to think that they are not answerable, yet no answer seems to be supportable by deductive or inductive reasoning. We could, of course, arbitrarily construct a principle from which an answer to any of these questions could be deduced--but the principle would be ad hoc, false, or, very likely, both ad hoc and false. Likewise, we could change these questions so that they could be answered inductively, by collecting evidence (e.g. Do a majority of English-speakers call tomatoes a kind of fruit or a kind of vegetable?)--but then we would be changing the questions into others. As they stand, these questions are questions about concepts, questions which deal with borderline examples and special cases, questions which have to do with what it means for something to be a vegetable, a lie, a murder, or a person. These are, as I have tried to show above, the sorts of questions judges face from time to time (Is this a patentable invention, or a mere advance, by mechanical steps, from previous inventions? Should the defendant have foreseen that there was a person lying on the road?). And they are the sorts of questions which generate moral conflicts, questions about the moral status of particular acts, questions about good and evil, virtue

and vice.

But enough has been said to be confident that the perplexity which may arise when the two features of these questions are considered together is caused by the dogma of Deductivism: these questions are, if, in the end, not always answerable, then always arguable; but they are not always answerable by either deductive or inductive means. Reasons can be advanced for and against, a case can be made one way or another. Still, if the perplexity caused by the dogma of Deductivism is removed, perplexity may still remain. These questions are perplexing. They are challenges to our ways of construing certain non-determinate concepts. They express the doubt that the concepts at issue are really as determinate as we may tend to take them to be.

It has been my claim throughout that, especially in philosophy, ethics and the law, the types of questions which are normally asked and which usually give rise to controversy and require judgement, are questions which require reflection. There is a systematic difference between formal logic, mathematics and the natural sciences, on the one hand, and law, ethics and philosophy on the other. And that difference, that systematic difference, has to do with the nature of the questions and problems which arise in systematic and determinate fields and those which arise in fields where optimal and universal determinateness is lacking. It is not possible to draw a hard and fast boundary between determinate and non-determinate fields, between, say, mathematics and philosophy, between mathematical questions and puzzles and philosophical questions and puzzles. And I am not suggesting that we should try. I am not interested in devising a Peircian classification of human knowledge--there are bound to be borderline cases of questions and

problems which do not fit neatly in any single branch of human knowledge. But I would like to suggest that there are paradigm cases of philosophical problems and paradigm cases of mathematical problems. And these cases can be identified if one has a familiarity with the respective fields.

But, if I am right and Deductivism is an avoidable dogma, what consequences does this fact have for the various fields where questions requiring reflection arise? Consider ethics. In the first chapter I sketched out some of the popular metaethical positions and tried to show how each seems to incorporate Deductivism. With regard to the question of the nature of moral reasoning, Deductivism tends to direct the discussion towards the general view that moral reasoning is reasoning from principles, that justification for any moral assertion requires an appeal to a moral principle. I would not, of course, want to claim that when we engage in moral argument we never reason from principles in a deductive manner. Rather, what I want to claim is that we have at our disposal a way of reasoning which is neither deductive nor inductive and which is suitable for moral questions.

This suggests a fairly modest metaethical thesis which I want to set out. The thesis is that moral assertions are arguable, that is, they are assertions for which the activity of adducing reasons for or against their truth or plausibility makes sense. On the face of it, this may not seem to be much of a thesis. It is not the thesis that, for example, moral assertions are claims which are objectively true or false; it is not even the thesis that they are claims capable of being true or false. It is rather the modest thesis that we are neither amazed or bemused, suspicious or incredulous when someone claims to be offering a reason in favour of a moral position, or proposes to argue for or against a particular moral judgement. We are neither amazed nor are we incredulous; rather

we find it very natural, and indeed quite proper to argue for or against a moral judgement or position. Moral assertions, then, may be viewed as utterances which count as moves (opening or otherwise) of possible arguments. They are assertions which invite some sort of argumentative response, supportive or not--a challenge for reasons for thinking it so, a request for clarification or modification, a rejoinder, a rejection, an affirmation, in short, an argument.

Usually, a moral issue is built out of one or more questions about one or more moral or (in the context) morally relevant concepts. We saw this in the case of the moral issue of abortion where the concepts of person, part of a body, right to life, right to one's own body and others joined to create the moral issue of whether it is morally permissible to have an abortion. Now, questions, problems or (even) puzzles about the nature of moral or morally relevant concepts account for why a moral issue is an issue at all. Admittedly, not everything which could be called a moral issue--not everything about which moral assertions are properly made--is conceptually perplexing. There are moral issues about the means one ought to follow to arrive at antecedently specified ends--issues about the best way of accomplishing what is agreed on all sides to be the best thing to accomplish given the circumstances. These too are moral issues which need to be resolved.

Utilitarians and rational decision theorists at times say things which would suggest that they think all moral issues are like this. For them, moral problems are usually and typically problems about choices, but choices in a specified and usually unquestioned matrix of choice--namely, maximising pleasure, social goods or some utility or other. The problem is, roughly, one of calculating, of adding up the pluses and

minuses in the various columns and choosing that course of action which maximises whatever it is that one ought to maximise. There are, undoubtedly, moral problems like this, problems which call for calculation. But these problems of calculation (some of which are extremely complicated) presuppose answers to other questions for which calculation is wholly inappropriate. (One cannot add up the pluses and minuses to determine whether the principle of utility is correct; one has to see if it is correct by looking at what follows from instances of its application.) These latter problems are problems about the moral nature of particular kinds of acts, problems about the ends, as well as questions about the goodness, badness, virtuousness and viciousness of particular things and acts. And these problems are problems about moral concepts. These are the issues which require reflection.

A moral argument might be thought of as a context in which the arguability of moral assertions can be manifested. Frequently, the problem which leads to moral conflict is a problem of determining what sort of thing or act the thing or act under consideration is. It is a problem parallel to the judge's problem of correctly describing the instant case from the point of view, the topos, of some area of the common law. Perplexing moral questions are difficult because reasons can be adduced both for and against one or another moral description. The difficult moral question concerns a case which has affinities with other cases, and these affinities take us into two different, perhaps two opposite directions. There is a point to be made comparing some cases of abortion with cases of taking an innocent life--exposing affinities which count as reasons for the immorality of abortion. Yet, there is also some point to be made comparing any case of abortion with

cases of allowing one to exercise one's right over one's own body.

Perplexing moral questions are typically those which are cited as examples of questions which are beyond the scope of reason, immune from or at least irresolvable by rational discussion. Here we may have, for example, conflicts over basic moral principles. These are the conflicts which persist and are, it is thought, inherently irresolvable. But the lessons of previous discussions suggest a different, more optimistic conclusion. Reflective reasoning is a remedy to the mistaken view that reasoning fails us in such cases of moral disputes where deduction and induction fail us. Moreover, if we come to see that persisting moral conflicts are evidence, not of the failure or built-in limitations of reasoning and argument, but of the difficulty of the conceptual issues involved, then by concentrating on the applicability of non-deductive, non-inductive reasoning in morality we shall be able to secure a confidence in the overall rationality of moral argument. For what, traditionally, are the views which are designed to shake this confidence?

One line of sceptical attack against the possibility of rational moral argument is founded on Hume's doctrine that there is a logical gap between moral assertions and statements of fact. The claim is that one cannot deduce moral assertions from statements of fact, but only from other moral assertions. Enough has been said to see why it is that Hume's doctrine is true; but more importantly, it should now be clear why the truth of Hume's doctrine has no bearing on the question of whether there can be rational and successful moral argument.

First, why is Hume's doctrine true? It is because moral concepts, like most legal concepts and like most of the concepts which interest

the philosopher, are not determinate. Radical, reductive naturalism is, if not false, then at least misleading as it is based on a misunderstanding of the nature of the rational connection between facts and moral concepts: moral concepts, it might be said, describe the world, but from the moral point of view (vide, Kovesi 146 ff.). Consider the legal parallel. If legal and legally relevant concepts like negligence, reasonable care, mens rea, right, foreseeability of harm, and the rest were always determinate from a legal point of view (if there were hard and fast rules, necessary and sufficient conditions for the application of these concepts), then it would be difficult to account for the deliberative role of judges and jurymen. Usually, the facts of a case are not sufficient to uniquely decide the question one way or another; the judge's or the jury's job is not to deduce from these facts a judgement. Their job is to reason to a decision; and that job requires that the deductive and inductive gap between the facts and the judgement be bridged, although not by a deductive or inductive inference.

But why shouldn't we be discouraged by this, why shouldn't we doubt that there can be a legitimate legal or moral argument? Just because we can and do reason over questions about non-determinate concepts. The facts of a legal case are relevant to a judgement, they count in favour of some judgements and against others; facts, in an argumentative context, count as reasons. In morality, the same is true. The hard-nosed empirical fact that aborted fetuses rarely survive to develop into adult human beings is relevant to the moral issue of abortion; that fact helps to make the question of abortion a morally significant one. If the facts were different, the moral issue would be different, or there might no longer be a moral issue. The fact

that aborted fetuses rarely survive is a reason--not a conclusive reason, but a reason--for thinking that there is something morally wrong about abortion. We have good conceptual reasons for thinking that there is a deductive gap between factual claims about abortion and moral assertions about it: Although abortion may be determinate from a factual point of view, it is not from a moral point of view; it is a morally significant act the moral assessment of which does not immediately follow from the kind of act it is. But that there is a deductive gap does not mean that there is a rational gap, an argumentative gap.

Another line of sceptical attack on the rationality of moral argument runs as follows: for a question to be answerable by argument, it must be answerable in a Yes or No fashion; either the question can be shown to admit of a definite answer one way or another, or else the question is not a genuine one. But there are moral questions--and perhaps the question of abortion is one of them--which appear not to be answerable in any definite way. So, they are not genuine questions, they are not questions amenable to rational discussion.

However, as consequences of Deductivism go, this one is the least plausible and is rarely defensible outside of very restricted contexts. For as we have seen even questions of fact cannot always be answered definitely one way or another. Is a mobile home a vehicle or a domicile? That depends on which one you are talking about (the one Aunt Elsie has been living in for the past twenty years, or the one Uncle Fred uses for his weekend fishing trips). Do the walls look green or blue? Well, in the early morning they do look blue, but then again when compared with the colour of the sofa they look green. Is this man bald? Well he is definitely going bald, but right now you can't really say (cf. Bambrough, 160 ff.). Or again: Is this forest land public or private

property? It's neither--it's public land on a ninety-nine year lease to a logging firm.

But, it might be further objected, even if we can successfully argue a moral assertion by presenting a persuasive case in its favour, still we cannot prove moral assertions. But this, I suggest, depends on what is meant by 'proof'. The Deductivist will claim that only deductive proof is proof. But this conception of proof (which is, in fact, a stipulative definition), while admirably suited to logic, mathematics and portions of the physical sciences, is really not appropriate here. The subject-matter of morality does not always admit of this sort of proof. To be sure, there can arise from time to time in an argumentative context conclusive, deductive disproof. If an argument in favour of a moral assertion relies upon or otherwise embodies an explicit contradiction, then the claim that it supports the moral assertion but not its negation can be, deductively, proved false. And showing someone that he has explicitly contradicted himself is pointing out to him that if he is to be rational he must give up one or the other or both of his contradictory premisses. But, for reasons mentioned above, it seems to me that explicit contradiction is a relatively rare phenomenon among minimally competent arguers. And while being shown that one has contradicted oneself ought to be an embarrassment, a setback, it is rarely a defeat--argument can continue once the contradictory claims are suitably modified to remove explicit contradiction.

But then, what sort of proof is appropriate for morality? The legal analogy can help to suggest a possible answer to this question. What kind of proof is a proof beyond reasonable doubt? First off, it is not necessarily deductively conclusive proof; it is an argument which

establishes a case to the satisfaction of a reasonable man without thereby removing the possibility of doubt. Legal proof is on the whole reflective proof. It is building a good case by means of a consideration of examples and parallels resulting in a decision about the instant case; and resulting as well in new conceptual understanding. Secondly, it is proof which breaks new ground, an argument which depends on our being able to show that the concepts and their instances which we are clear about are relevant to other concepts and their instances which we may not be clear about. It is thus an exploratory proof, a proof which sets out a new picture of what it means to be negligent, a murderer or a thief which stands as the correct picture until it is improved upon when new cases, new permutations of the possible ways one can be negligent, a murderer or a thief come before the court. At some point in the development of the tort of negligence, for example, it had to be argued that statutes which impose liability for negligent acts ought to be extended to cover negligent statements. At that point the topos of the argument consisted of, inter alia, hundreds of cases of negligent acts, instances of the application of the concept negligence. The reasoning which was required was reasoning about the relevant similarity, in so far as negligence as a tort was concerned, of acts and statements. When considering the possibility of negligent statements, judges had to look again at the picture of what it means for someone to be negligent. Before, acting negligently resulted in some sort of physical injury; now being negligent might not involve physical injury at all, but rather pecuniary damage. Can the concept survive this modification? But if it cannot then what about the affinity between the careless act and the careless word? And oughtn't one be liable for damages if the damage was

the consequence of not foreseeing the harm which a careless remark could make? These were the sorts of questions asked, and answered.

But shouldn't we say that legal proof is inconclusive? It is not always deductively conclusive; but it may be, for all of that, conclusive on its own terms. It should be noted that the conclusiveness which deductive proof displays is conclusiveness bought at a high price: deduction is demonstrative because it is non-ampliative. With legal proof, as with moral proof, we are successful if the case we make can withstand counter-argument, if whatever doubt expressed about its adequacy is unreasonable or unsupportable, and if counter-arguments are open to obvious objections. Here is a sort of conclusiveness: an argument can be conclusive if the price of standing by a counter-argument is so high that one finds oneself forced to abjure from rational discussion.

Now as we saw in Chapter II, the price of standing by the claim that a valid deductive inference is illegitimate is self-contradiction, at some point in one's defence of that claim one must contradict oneself. To be sure, this feature of deductive validity in large part rests on the strict characterisation of particles like 'not' and 'and' and 'if... then___' which the logician favours. But it is nonetheless a good argument for the conclusiveness of valid deductive proof that one finds oneself in such an awkward position when one tries to argue against the legitimacy of the proof. But, notice, this argument is not itself a deductively conclusive one. Moreover, in everyday argumentative contexts, in the courtroom and in rarified philosophical contexts we are perfectly willing to grant that there are convincing arguments, that there are proofs beyond reasonable doubt. But it is surely not the

case for all of these that our recognition that they are convincing and beyond reasonable doubt is the recognition that for any counter-argument one could devise there would be lurking self-contradiction. Not every proof is conclusive in the way that deductive proofs are, but some of these are still conclusive.

Suppose I hear little Johnny practising his piano every morning at seven. I hear him telling his friends that he can't come out and play with them because practising is more fun. And every conversation I have had with him has shifted to the topic of music. Moreover, he avidly reads stories about famous pianists, saves up his allowance money to buy sheet-music and plays his piano with an intensity and feeling that signals a budding virtuoso. Now, if I add all this up the result is a good case, a convincing proof that Johnny, unlike many his age, does not hate the piano. But, of course, he might. All of my evidence is compatible with his having a domineering father who threatens his son with corporal punishment if he doesn't pretend in every circumstance to love music and his piano. Fair enough; that is a possibility, but it is not very likely. All this time, he was pretending? He's so young to be such a convincing actor. Wouldn't he have slipped just once if he had been pretending all along? The very young have a difficult time covering stress; how could Johnny have hid the fact that his father was threatening him? And on and on. It is possible, but not a likely hypothesis. But now, if the only proof I could give for Johnny's not hating the piano was an argument the denial of which would involve self-contradiction, then I could never prove that he did not hate the piano. And that is surely suspicious.

It is suspicious because the standard for what is to count as

proof in a case like this is out of reach. And this is the sceptic's favourite trick. But he must be careful, for if the standard for proof is raised high enough, then nothing, not even deductive proof, will count as proof. And then the sceptic could not make his own case. A Cure for this sort of scepticism may very well be a strong dose of Pragmatism. Reasoning, as Peirce saw, has a purpose; we have an interest in reasoning properly because the success of our reasoning may have a direct effect on the success of our lives. But if success in reasoning is nothing short of deductive conclusiveness, then our chances of succeeding in the vast majority of our arguments are very slim indeed. If it weren't the case that we normally reason non-deductively, then for survival's sake if for no other we would have to invent non-deductive reasoning. The sceptic is well aware of this, he like the rest of us has a fairly good idea of when proofs are beyond reasonable doubt and when they are not. He couldn't put forward his own case unless he had some idea of what a good case for scepticism would be like.

So, if proof beyond reasonable doubt is not deductively conclusive, it is still conclusive. And without proof beyond reasonable doubt, deductive proof would not be possible, since the basis for deductive conclusiveness is itself arguable, though not deductively arguable. At some stage we must accept proof beyond reasonable doubt for the necessity of an axiom or the legitimacy of a rule of inference. The concept of proof can easily accomodate both deductive, conclusive proof and nondemonstrative proof beyond reasonable doubt; indeed, without the former it would not be the familiar concept and without the latter it would be wholly vacuous.

There is obviously much more that could be said about the nature and rationality of moral, and legal, argumentation. Once Deductivism

in ethics is abandoned and the deductive model of justification and proof is replaced by the reflective model, then the door is open to accounts of moral reasoning which do not rely on moral principles and reasoning from principles. And it is my belief that these accounts, once fully developed, may count as presumptions in favour of the objectivity of morality. For if we set aside the view that the only legitimate and respectable ways in which to reason over moral questions are deductive and inductive ways of reasoning, then we are at once liberated from the constraints which deduction and induction place on the moral sphere. We shall no longer feel obliged to force moral reasoning into a deductive mould and be disappointed by the fit. Reasoning from examples and parallels is, on the other hand, eminently suitable to the sorts of issues which are involved in moral disputes. It is because reflective reasoning is a rational way of exploring the nature of the concepts which give rise to moral questions. Moral issues--and to a lesser extent, legal issues--often strike us as being irresolvable because either there are no easy and immediate answers forthcoming--and the labour required to reach an answer seems either beyond human abilities or else Sisyphean--or else we can see the plausibility of both sides of the issue. But, if what I have been urging is correct, the difficulty and complexity of some moral issues is a consequence of the fact that these issues turn on how certain concepts are to be understood, and these concepts are non-determinate. But reasoning about non-determinate concepts is possible: reflection is not a calculus designed to solve moral issues in a unique way; but it is a way of finding one's way around, of proceeding rationally in a difficult and complex area of human knowledge.

The sceptic in morality and in the law is, of course, eager to show that the open-textured and 'fuzzy'-boundaried concepts which are the loci of debate are, because non-determinate, indeterminate, and hence that, in the end, there is no rational way of deciding when a particular case falls under one rather than another description. The sceptic tries to argue that since we don't have strict definitions, since we lack necessary and sufficient conditions for the application of concepts like person or negligent act, we have no clear idea of when something is a person or when an act is negligent. But this is a trap which can be avoided. To say that we are not always clear about the application of a non-determinate concept is not to say that we never are. And the value of reasoning from particulars, from examples and parallels, lies in the fact that we can argue from cases of the application of a concept which are clear to those which are less clear. The question of whether a human fetus is a person can be argued by exploring the affinities between fetuses and clear cases of persons, by exploring the affinities between fetuses and clear cases of non-persons, by comparing and contrasting. Questions about concepts stand in need of reasoning about concepts; and such reasoning, in the end, is always reasoning from case to case.

But what about philosophical questions? These are paradigm examples of questions about concepts. Moreover, they are often questions which arise out of strange-sounding, even paradoxical claims. John Wisdom, in his unpublished lectures Proof and Explanation and in several published articles, has sought to bring to our attention the nature of philosophical (or what he calls 'metaphysical') questions and disputes. His concern throughout is methodological, it is what he calls the

'metaphysics of metaphysics'. He is interested in showing why strange-sounding, even paradoxical philosophical questions are genuine questions, and why reasoning from case to case is the appropriate tool for philosophical reasoning. I want to conclude by saying a bit more about Wisdom's conception of metaphysical questions and philosophical methodology as this relates to what has been said above.

Wisdom is fond of pointing to philosophical questions which sound absurd or paradoxical. Some of the examples he gives are examples of questions which have an initial fascination because they ask whether a particular, familiar concept applies in cases where we may have some doubt whether it applies--Do dogs think? Do babies think? These are akin to some of the questions which I have offered as questions requiring reflection in ethics and the law. One might suppose that the question whether babies think has immediate moral and legal significance. And the question asked of dogs may very well have moral significance: if they think, do they have rights? Are they sentient? Do they command our moral respect? As it happens, the law has already delved into such questions in the process of ruling on the application of cruelty to animals legislation.

But Wisdom has other sorts of questions in mind. In particular he considers questions which develop out of questions about borderline cases. These are questions which push borderline case questions, as we might say, to the end of their tether--Does anyone else besides me think? There is a distinct change in character in evidence here, and Wisdom is keen to bring it to our attention. The questions about dogs and babies ask whether the concept of a thinking being can be correctly applied to dubious or borderline cases--granting that the concept has

clear instances. But this last question--the full-blown, metaphysical question--asks of the clear instances whether they apply. The question is peculiar, paradoxical: it asks whether we really know what no one would doubt we know; it pushes us to consider what is obvious in the same spirit as we consider what is not obvious. And it is crucially important that we take this sort of question seriously: in the course of addressing ourselves to the absurd suggestion that there is some doubt that others think, we may come to understand why that doubt is excessive; and out of that may come a clearer understanding of what, of course, we know already.

Wisdom frequently notes an important feature of the full-blown metaphysical question. With these questions it is vain to try to deduce the answer we want, since to do so one must employ premisses which claim, in effect, that part of what is at issue is not at issue. For example, consider the metaphysical query, Do we know anything about the past? Suppose we decide to deduce the answer that we do indeed have knowledge about the past. Suppose we set out the following major premiss: If one remembers that P, then one knows something about the past, namely that P. This premiss, in conjunction with another--say, I remember winding my watch this morning--would yield the conclusion that I know something about the past. But this is, obviously enough, no resolution of the metaphysical worry. It is not because (i) the transition from 'knows something about the past' to 'remembers that P' leaves the problem untouched: Do we ever really remember anything? is the same worry we began with; and (ii) the minor premiss is a claim about the past, it is a claim about what precisely is at issue. In short: nothing from which the paradoxical (or, for that matter, the

borderline) is deducible will not itself be paradoxical (or borderline).

Another important aspect of metaphysical questions of the sort Wisdom is interested in is that the question does not wait upon new information; it is a question which remains after, as Wisdom puts it at one point, 'the whole tapestry has been unrolled' ((3)). If one is considering the metaphysical problem of the existence of other minds, the issue is not one which waits on new facts being brought to light. There is no need to conduct surveys, nor is there any need to watch people very closely to see if new behaviour, previously undetected, can be noted which would cinch it. The facts are all there before us: and that is why the question is both odd and perplexing.

Now, both of these features of Wisdom's metaphysical questions are features of the questions considered above in ethics and the law. Deduction is often ineffectual as a means of discovering the solution to a problem about a moral or a legal concept. And for the same reason. We could say that negligence is just an unintentional failure to foresee harm; but this translation need not aid the judge in determining whether the case before him is a case of negligence: Whatever reflection is required to see the instant case as one of negligence is just the reflection required to see it as a case of unintentional failure to foresee harm. Moreover, with legal questions there is a natural limit to the factual information relevant to a case which can be collected. When all that can be brought out has been, the counsel for the defendant and the counsel for the plaintiff (or the prosecutor) try to fit the facts into different patterns, patterns which have legal significance, patterns which point either to the judgement of guilt or to innocence. The job of the judge or the jury is to see if the opposing descriptions

fit the facts, if the relevant legal and legally significant concepts are being applied in a manner which would satisfy the rational man.

At times, Wisdom's examples of metaphysical questions tend to strike one as being, at least in the context of the everyday work in the profession, a bit grandiose: not all philosophical problems have a ring of paradox to them, not all philosophical worries have that hint of undefinable strangeness one finds with Do other minds exist? or Could we all be mad? or Can we ever know anything about the past? Wisdom is, of course, bringing our attention to the paradigm examples of metaphysical questions and claims for a reason: he wants us to reject the view that such questions are just absurd, just unconfirmable by further investigation, or just a matter of words. Still, philosophers are more likely to be concerned with 'fine-tuning' philosophising, and are more likely to concentrate their efforts on tamer questions: Do all verbs of perception have intentional uses? Can Jones believe that the shortest spy is a spy if there is no particular man whom Jones believes to be a spy? Could a person be through and through rational yet amoral? Are pains perceptual objects, like sounds? Is a mental particular like a feeling of well-being something which has temporal position and extension but not spatial position and extension? These questions (gleaned from randomly selected pages of a current issue of Mind) lack the air of paradox, though they would doubtless sound peculiar to the ear of a non-philosopher. Yet, there are questions about concepts; paradoxical or not, they are questions which require reflection.

One last feature of metaphysical questions which Wisdom brings out is important here. Metaphysical questions are often questions about concepts which are part of a non-philosophical theory, say physics,

psychology, morality, the law or theology. But, as Wisdom (after Wittgenstein) insists, metaphysical disputes over the notions of time, simultaneity, the unconscious, the good, mens rea, grace or whatever are not disputes which take place within these theories and fields; the philosopher does not contribute to physics, psychology, morality, law or theology, he is engaged in meta-criticism. The philosopher is interested in the disputes and controversies in fields other than his own; he is concerned to understand the explanations, justifications and proofs in other fields. Yet, philosophical criticism does not add to the body of findings or theories of other fields; rather, philosophical questions are as they are because the findings and theories are as they are.

This suggests that the topos of a philosophical dispute may be determined in part by what is--at the time of the dispute--accepted physics or current psychology, what is accepted moral practice or the present stage in the development of the law. Put in a slightly different way, the topos of a philosophical dispute over some concept which has a place in, say, physics, contains the practice of physics. The practice is given; philosophical criticism neither adds to nor subtracts from the practice as it is. What philosophical criticism is aimed at is a fuller understanding of the logical structure of the practice, as mirrored in the concepts under consideration. As the object of philosophical investigation changes, so too does the argumentative topos; part of this topos is thus, in a sense, inviolable: philosophical criticism is an attempt to provide a clear picture of a discipline like physics or psychology or law without thereby contributing to the subject-matter. The rest of the philosophical topos is itself always

open to philosophical criticism--when the need is felt for increased conceptual understanding of the presuppositions of the discussion.

In this thesis, an attempt has been made to engage in reasoning about reasoning. Without attempting to contribute to the two areas from which most of the examples have been drawn--namely, morality and the common law--I have endeavoured to set out the practice of reasoning in these fields, to get an idea of the picture of rationality in these fields, and, at least in part, to understand the nature of justification and of proof in these non-determinate fields. My emphasis throughout has been on the nature of a certain class of questions--set in argumentative contexts provided by the field in question--which require conceptual understanding. The concepts about which I have been concerned--the concepts of reasoning, rationality, justification and proof--themselves generate questions which require reflection. It is my hope that what has gone before counts as a preliminary attempt to become clear about these crucially important concepts.

BIBLIOGRAPHY

- Allen, C. K. Law in the Making. London: Oxford U.P., 1927.
- Aristotle (1) The Works of Aristotle. Edited by W. D. Ross. Vol. I.
London: Oxford U.P., 1953.
- _____. (2) The Works of Aristotle. Vol. VIII.
- _____. (3) The Works of Aristotle. Vol. IX.
- _____. (4) The Works of Aristotle. Vol. XI.
- Ayer, A. J. Language, Truth and Logic. New York: Dover Publishing Inc., 1946.
- Ayers, M. R. "Counterfactual and Subjunctive Conditionals." Mind LXXIV (1965): 347-364.
- Bambrough, J. R. "Unanswerable Questions." Proceedings of the Aristotelian Society, Supplementary Volume XL (1966): 151-172.
- Beardsmore, R. W. Moral Reasoning. London: Routledge and Kegan Paul, 1969.
- Becker, L. C. "Analogy in Legal Reasoning." Ethics LXXXIII (1973): 248-255.
- Bickenbach, J. E. "The Status of the Propositions of the Tractatus." Dialogue XIII (1974): 763-772.
- Blackburn, S. Reason and Prediction. Cambridge: Cambridge U.P., 1973.
- Braithwaite, R. B. Scientific Explanation. Cambridge: Cambridge U.P., 1953.
- Cardozo, B. N. The Nature of the Judicial Process. New Haven: Yale U.P. 1965.
- Carroll, L. "What the Tortoise Said to Achilles." Mind, N.S., IV (1895): 278-280.
- Cicero. Topica, in De Inventione; De Optimo Genere Oratorum; Topica. Translated by H. M. Hubbell. London: W. Heinemann Ltd., 1959.
- Cohen, L. C. The Implications of Induction. London: Methuen and Co., Ltd., 1970.

- Cohen, M. R. and Cohen, F. S., editors. Readings in Jurisprudence and Legal Philosophy. New York: Prentice-Hall, Inc., 1951.
- Cross, R. Precedent in English Law. Oxford: Oxford, Clarendon Press, 1968.
- Denning, Lord. The Road to Justice. London: Stevens and Sons, 1955.
- Dias, R. W. M. Jurisprudence. 2nd ed. London: Butterworth, 1964.
- Dummett, M. A. E. "The Justification of Deduction." Proceedings of the British Academy. Volume LIX. Annual Philosophical Lecture. London: Oxford U.P., 1973.
- Edel, A. (1) Ethical Judgement: The Use of Science in Ethics. London: Free Press of Glencoe, 1955.
- _____. (2) Method in Ethical Theory. London: Routledge and Kegan Paul, 1963.
- Edgley, R. Reason in Theory and Practice. London: Hutchinson, 1969.
- Ewing, A. C. Ethics. London: Teach Yourself Books, St. Paul's House, 1973.
- Feigl, H. "De Principiis Non Disputandum...?" Reprinted in Philosophical Analysis, M. Black, editor. Englewood Cliffs, N.J.: Prentice-Hall, (1963): 113-147.
- Findlay, J. N. Language, Mind and Value. London: Allen and Unwin, 1963.
- Flew, A. Thinking About Thinking. Glasgow: Fontana, 1975.
- Frank, J. Law and the Modern Mind. Garden City, N.Y.: Doubleday Co., 1963.
- Gewirth, A. "Positive "Ethics" and Normative "Science"." Reprinted in (84), 27-47.
- Goodhard, A. L. "Determining the Ratio Decidendi of a Case". 40 Yale Law Journal (1930), 161.
- Goodman, N. Fact, Fiction and Forecast. Indianapolis: Bobbs-Merrill, 1965.
- Haack, S. "The Justification of Deduction." Mind LXXXV (1976): 112-119.
- Hamblin, C. L. Fallacies. London: Methuen and Co., 1970.
- Hampshire, S. "Fallacies in Moral Philosophy." Reprinted in (53), 157-176.

- Hare, R. M. (1) The Language of Morals. Oxford: Oxford U.P., 1952.
- _____. (2) Freedom and Reason. Oxford: Oxford U.P., 1963.
- Hart, H. L. A. The Concept of Law. Oxford: Clarendon Press, 1975.
- Henkin, L. "'Neutral Principles' and Future Cases" in (43), 301-309.
- Holmes, O. W. (1) "The Path of the Law." 10 Harvard Law Review (1899): 457.
- _____. (2) The Common Law. Boston: Little Brown and Co., 1923.
- Hook, S., editor. Law and Philosophy. New York: New York U.P., 1964.
- Hume, D. Enquiry Concerning the Principles of Morals. Edited by L. A. Selby-Bigge. Oxford: Oxford, Clarendon Press, 1961.
- Hutcheson, J. C., Jr. "The Judgment Intuitive: The Function of the 'Hunch' in Judicial Decision." Reprinted in (49), 407-419.
- Jones, H. W. "The Practice of Justice." Reprinted in (49), 425-436.
- Kant, I. (1) Fundamental Principles of the Metaphysic of Ethics. Translated by T. K. Abbott. London: Longman, 1969.
- _____. (2) Prolegomena to Any Future Metaphysics. New York: Bobbs-Merrill, 1960.
- Kent, E. A., editor. Law and Philosophy. New York: Appleton-Century-Crofts, 1970.
- Kovesi, J. Moral Notions. New York: Routledge and Kegan Paul, 1971.
- Levi, E. H. An Introduction to Legal Reasoning. Chicago: University of Chicago Press, 1972.
- Lloyd, D. The Idea of Law. Middlesex: Penquin Books, 1964.
- Margolis, J. (1), editor. Contemporary Ethical Theory. New York: Random House, 1966.
- _____. (2) "Parallels in Moral and Aesthetic Discourse." Journal of Aesthetic Education I (1968-1969): 65-77.
- Mill, J. S. A System of Logic. Volumes I and II. London: Longmans, Green and Co., 1965.
- Montefiore, A. A Modern Introduction to Moral Philosophy. London: Routledge and Kegan Paul, 1958.
- Murphy, J. G. "Law Logic." Ethics LXXVII (1967): 193-201.

- Newell, R. W. The Concept of Philosophy. London: Methuen, 1967.
- Olson, R. G. Meaning and Argument. New York: Harcourt, Brace and World, 1969.
- Paton, G. W. A Textbook of Jurisprudence. Oxford: Oxford, Clarendon Press, 1946.
- Peirce, C. S. Collected Papers. Edited by Arthur Burks. Cambridge, Mass.: Harvard U.P., 1958.
- Perelman, Ch. (1) The Idea of Justice and the Problem of Argument. London: Routledge and Kegan Paul, 1963.
- _____. (2) "Rhetoric and Philosophy." Philosophy and Rhetoric I (1968): 15-24.
- Perelman, Ch. and Olbrechts-Tyteca, L. The New Rhetoric: A Treatise on Argumentation. Notre Dame, Illinois: University of Notre Dame Press, 1969.
- Perry, T. D. Moral Reasoning and Truth. Oxford: Oxford, Clarendon Press, 1976.
- Phillips, D. Z. and Mounce, H. O. Moral Practices. New York: Schocken Books, 1970.
- Plato. The Dialogues of Plato. Translated by B. Jowett, Vols. I and II. New York: Random House, 1937.
- Popper, K. R. Conjectures and Refutations. New York: Harper and Row, 1963.
- Pound, R. "Mechanical Jurisprudence." 8 Colorado Law Review (1908).
- Prichard, H. A. Moral Obligation. Oxford: Oxford U.P., 1971.
- Quine, W. v. O. Methods of Logic. London: Routledge and Kegan Paul, 1962.
- Rawls, J. A Theory of Justice. Cambridge, Mass.: Harvard U.P., 1971.
- Rescher, N. (1) The Coherence Theory of Truth. London: Oxford, 1973.
- _____. (2) "Foundationalism, Coherentism, and the Idea of Cognitive Systematization." Journal of Philosophy LXXI (1974): 695-708.
- Ross, W. D. (1) The Right and the Good. Oxford: Oxford U.P., 1973.
- _____. (2) Foundations of Ethics. Oxford: Oxford, Clarendon Press, 1968.

- Russell, B. The Problems of Philosophy. London: Thornton Butterworth Ltd., 1912.
- Salmon, W. "The Concept of Inductive Evidence." Reprinted in (83), 48-57.
- Sellars, W. "Are There Non-Deductive Logics?" Reprinted in W. Sellars Essays in Philosophy and Its History. Dordrecht: Reidel Publishing Co., (1974): 417-430.
- Sextus Empiricus. Works, 4 vols. Translated by R. G. Bury. London: Loeb Classical Library, 1935-1949.
- Stone, J. Legal System and Lawyers' Reasoning. London: Stevens and Sons, 1964.
- Strawson, P. F. An Introduction to Logical Theory. London: Methuen and Co., 1952.
- Swinburne, R., editor. The Justification of Induction. Oxford: Oxford U.P., 1974.
- Thomson, J. J. and Dworkin, G., editors. Ethics. New York: Harper and Row, 1968.
- Thomson, J. J. (1) "Reasons and Reasoning." Edited by Max Black. Philosophy in America. London: G. Allen and Unwin, Ltd., (1965): 282-303.
- _____. (2) "A Defense of Abortion." Philosophy and Public Affairs I (1971): 47-66.
- Toulmin, S. The Uses of Argument. Cambridge: Cambridge U.P., 1969.
- Urmson, J. O. "Some Questions Concerning Validity." Reprinted in (83), 74-84.
- Waismann, F. How I See Philosophy. Edited by R. Harre. New York: Macmillan, St. Martin's Press, 1968.
- Wiener, P. P., editor. Charles S. Peirce: Selected Writings. New York: Dover, 1958.
- Williams, G. "Language and the Law." Reprinted in (18), 577-580.
- Wisdom, J. (1) Philosophy and Psychoanalysis. Oxford: Basil Blackwell, 1953.
- _____. (2) Paradox and Discovery. Oxford: Basil Blackwell, 1965.
- _____. (3) "Proof and Explanation--The Virginia Lectures." Unpublished notes of lectures given at the University of Virginia, 1957.

Wittgenstein, L. (1) Tractatus Logico-Philosophicus. Translated by D. Pears and B. F. McGuinness. London: Routledge and Kegan Paul, 1968.

_____. (2) The Blue and Brown Books. Oxford: Basil Blackwell, 1958.

_____. (3) Philosophical Investigations. Oxford: Basil Blackwell, 1958.

_____. (4) On Certainty. Oxford: Basil Blackwell, 1964.

von Wright, G. H. The Logical Problem of Induction. Oxford: Basil Blackwell, 1965.

REQUEST FOR DUPLICATION

I wish a photocopy of the thesis by

J. Bickon back (author)

entitled The Nature + Structure of Analogical Reasoning

The copy is for the sole purpose of private scholarly or scientific study and research. I will not reproduce, sell or distribute the copy I request, and I will not copy any substantial part of it in my own work without permission of the copyright owner. I understand that the Library performs the service of copying at my request, and I assume all copyright responsibility for the item requested.

B30191